



US009052830B1

(12) **United States Patent**
Marshak et al.

(10) **Patent No.:** **US 9,052,830 B1**
(45) **Date of Patent:** **Jun. 9, 2015**

(54) **TECHNIQUES FOR AUTOMATED EVALUATION AND MOVEMENT OF DATA BETWEEN STORAGE TIERS FOR THIN DEVICES**

(75) Inventors: **Marik Marshak**, Newton, MA (US); **Adnan Sahin**, Needham, MA (US); **Xiaomei Liu**, Natick, MA (US); **Hui Wang**, Upton, MA (US); **Sean C. Dolan**, Belmont, MA (US); **Amnon Naamad**, Brookline, MA (US)

(73) Assignee: **EMC Corporation**, Hopkinton, MA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 912 days.

(21) Appl. No.: **13/135,306**

(22) Filed: **Jun. 30, 2011**

(51) **Int. Cl.**
G06F 12/00 (2006.01)
G06F 13/00 (2006.01)
G06F 13/28 (2006.01)
G06F 3/06 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 3/064** (2013.01); **G06F 3/0611** (2013.01); **G06F 3/0608** (2013.01); **G06F 3/067** (2013.01)

(58) **Field of Classification Search**
CPC . G06F 12/0223; G06F 9/4406; G06F 3/0611; G06F 3/0644; G06F 3/0647; G06F 3/0665; G06F 3/0685
USPC 711/202, E12.078, 118
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,124,272 B1 * 10/2006 Kennedy et al. 711/173
7,721,044 B1 * 5/2010 Chatterjee et al. 711/114

7,949,637 B1 * 5/2011 Burke 707/655
7,953,929 B1 * 5/2011 Chatterjee et al. 711/114
8,433,848 B1 * 4/2013 Naamad et al. 711/114
8,566,546 B1 * 10/2013 Marshak et al. 711/165
8,566,553 B1 * 10/2013 Marshak et al. 711/170
8,572,338 B1 * 10/2013 Marathe et al. 711/162
8,583,838 B1 * 11/2013 Marshak et al. 710/18
8,935,493 B1 * 1/2015 Dolan et al. 711/161
2009/0070541 A1 3/2009 Yochai
2009/0150593 A1 * 6/2009 Hamilton et al. 711/101
2011/0246740 A1 * 10/2011 Yata et al. 711/165
2012/0005435 A1 * 1/2012 Emaru et al. 711/154
2012/0131196 A1 * 5/2012 Yamada 709/226
2013/0145092 A1 * 6/2013 Miwa et al. 711/114

OTHER PUBLICATIONS

U.S. Appl. No. 12/924,361, filed Sep. 24, 2010, Marshak, et al.
U.S. Appl. No. 12/924,396, filed Sep. 25, 2010, Marshak, et al.
U.S. Appl. No. 12/798,097, filed Mar. 30, 2010, Marshak, et al.
U.S. Appl. No. 12/803,571, filed Jun. 30, 2010, Marshak, et al.

(Continued)

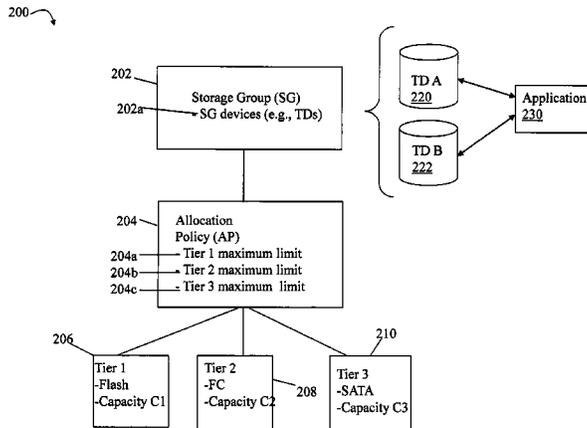
Primary Examiner — Sheng-Jen Tsai

(74) Attorney, Agent, or Firm — Muirhead and Saturnelli, LLC

(57) **ABSTRACT**

Described are techniques for evaluating data movement alternative. A set of criteria including capacity and performance limits is received. First processing is performed to evaluate a plurality of alternatives for use in data movement with respect to a set of logical devices having data stored on a set of physical storage devices. Each of the plurality of alternatives includes a different set of data movement criteria comprising capacity limits and a different set of performance limits. The set of physical storage devices includes at least a first physical device of one of a plurality of storage tiers and a second physical device of another one of the plurality of storage tier. One of the sets of performance limits is selected in accordance with the first processing.

20 Claims, 22 Drawing Sheets



(56)

References Cited

OTHER PUBLICATIONS

U.S. Appl. No. 12/803,570, filed Jun. 30, 2010, Marshak, et al.
U.S. Appl. No. 12/924,396, filed Sep. 27, 2010, Marshak, et al.

U.S. Appl. No. 12/924,430, filed Sep. 27, 2010, Marshak, et al.
U.S. Appl. No. 13/135,261, filed Jun. 30, 2011, Marshak, et al.
U.S. Appl. No. 13/135,265, filed Jun. 30, 2011, Sahin, et al.

* cited by examiner

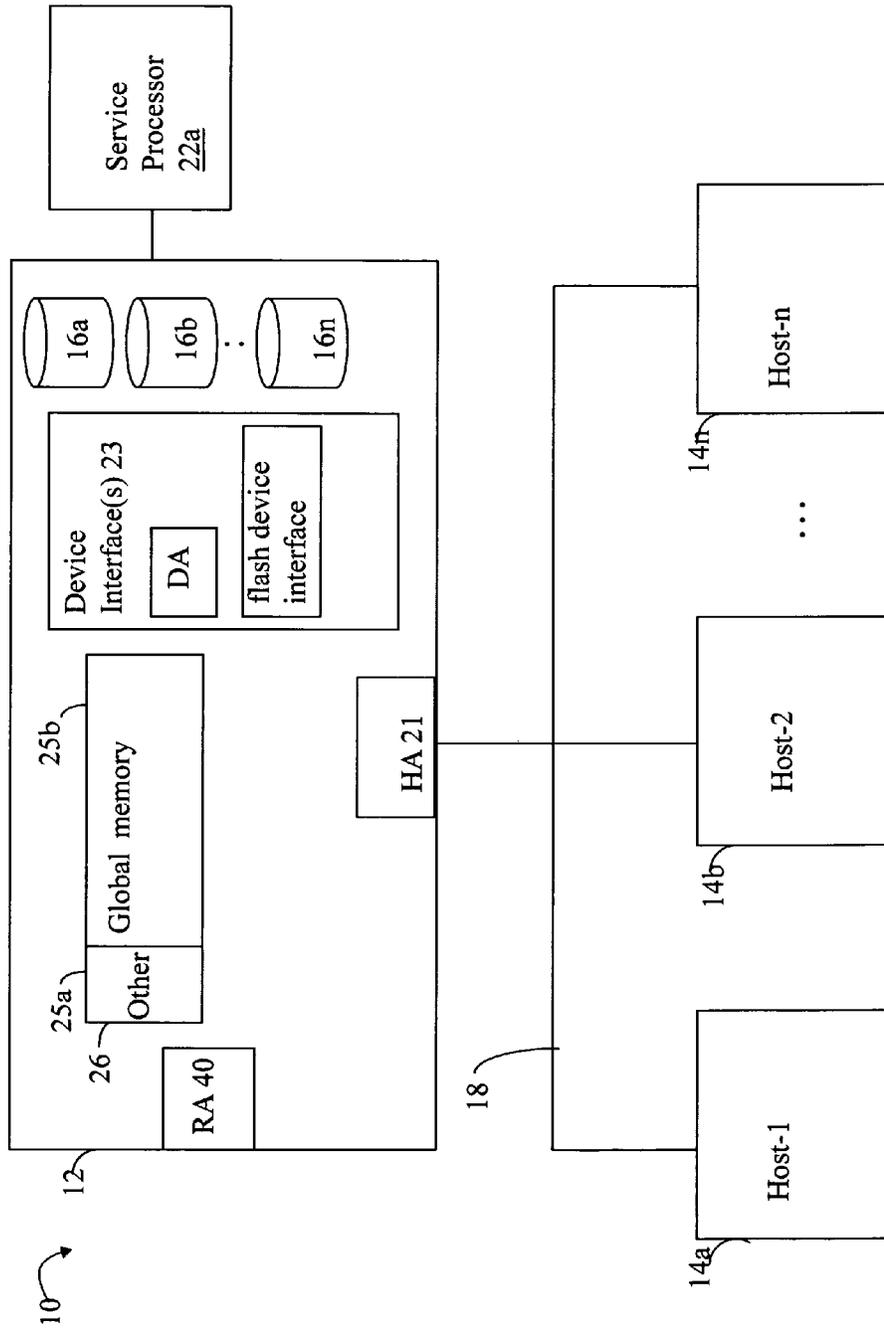


FIG. 1

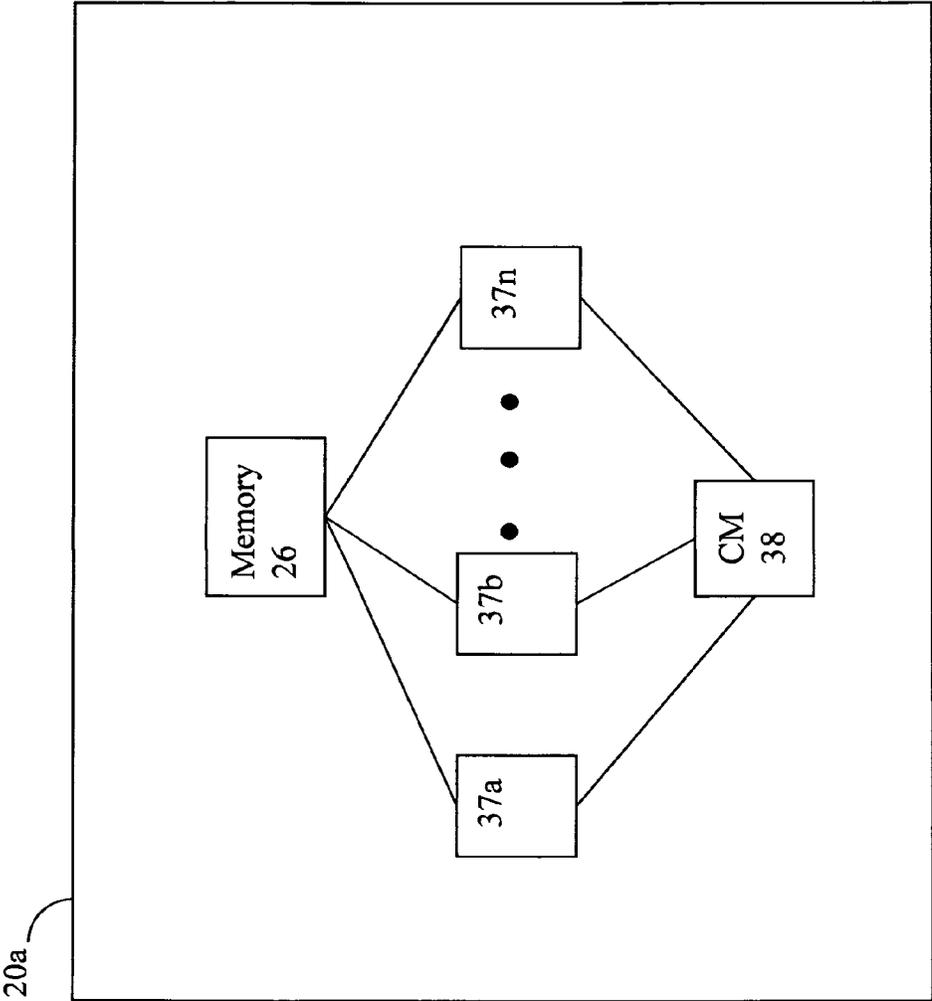


FIG. 2

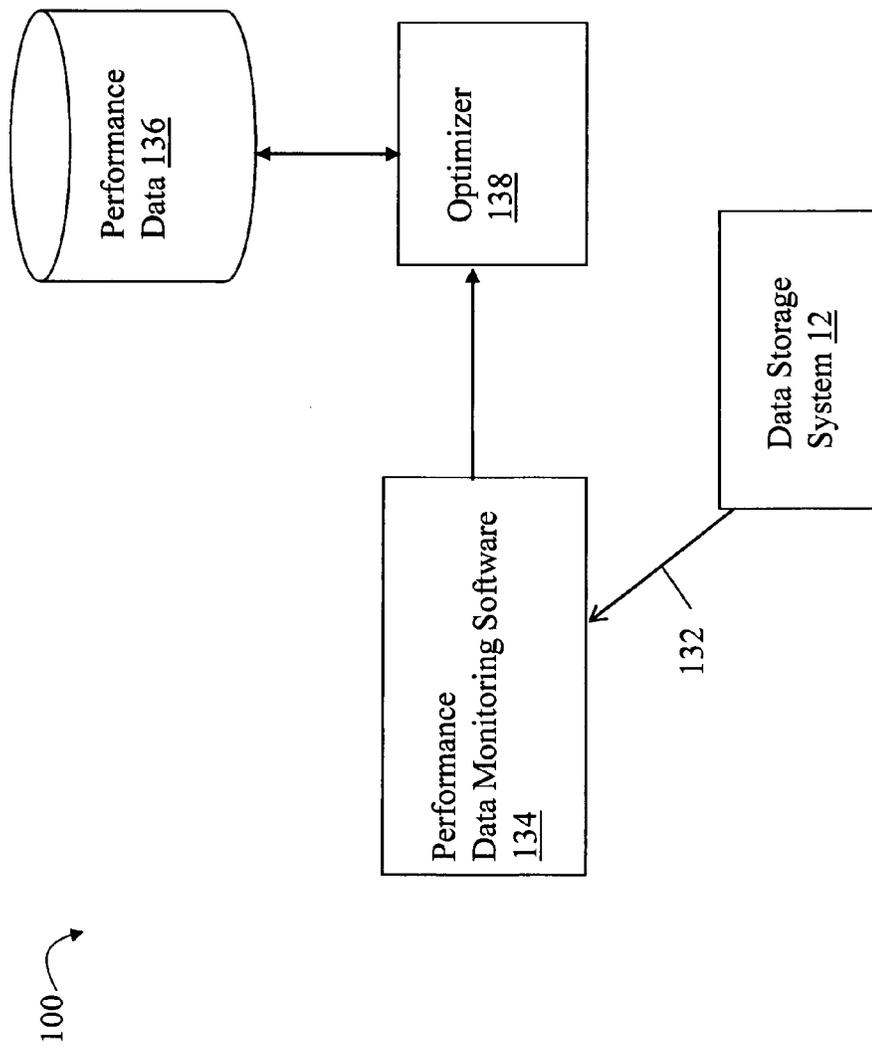


FIG. 3

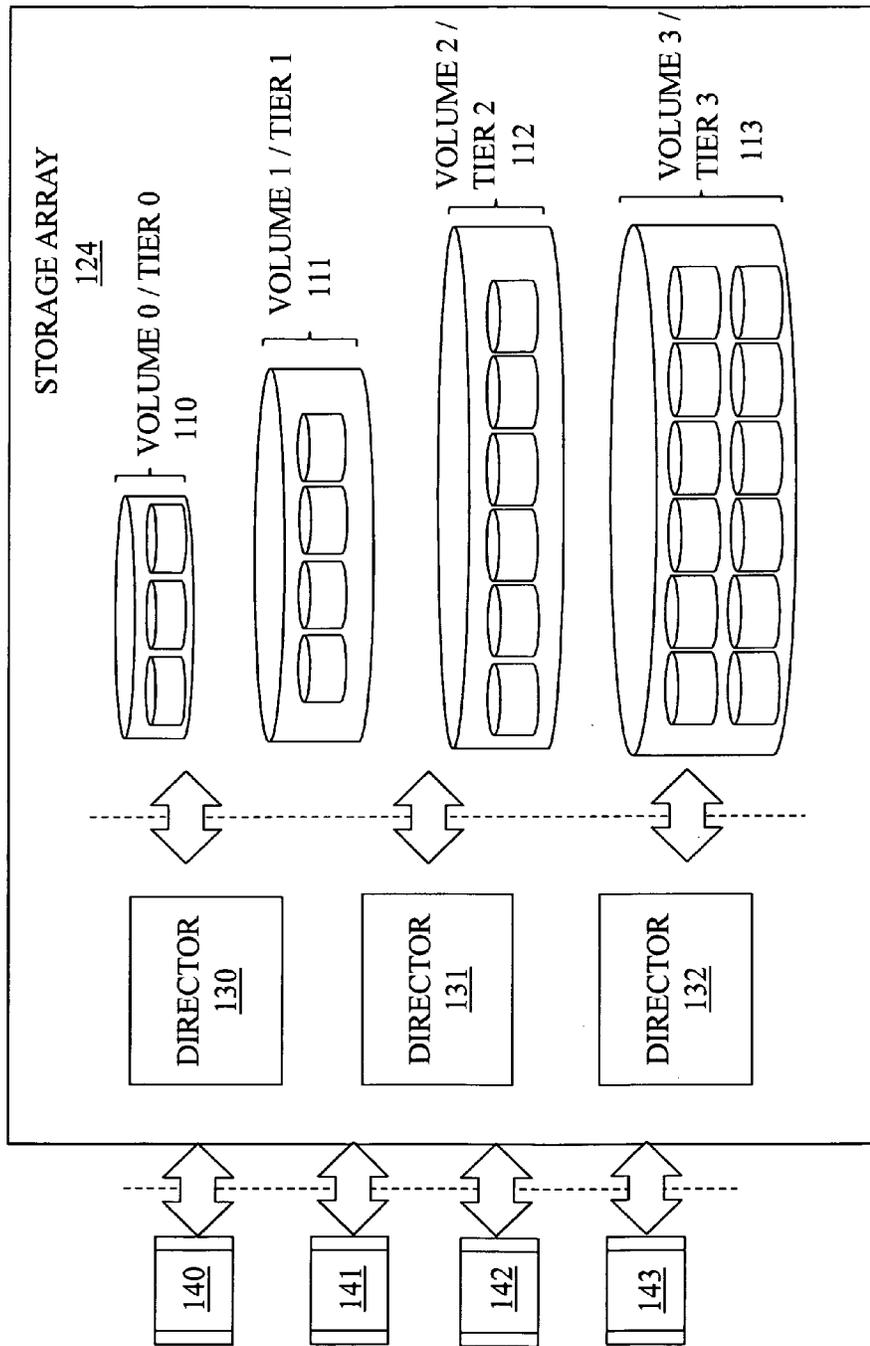


FIG. 4 150

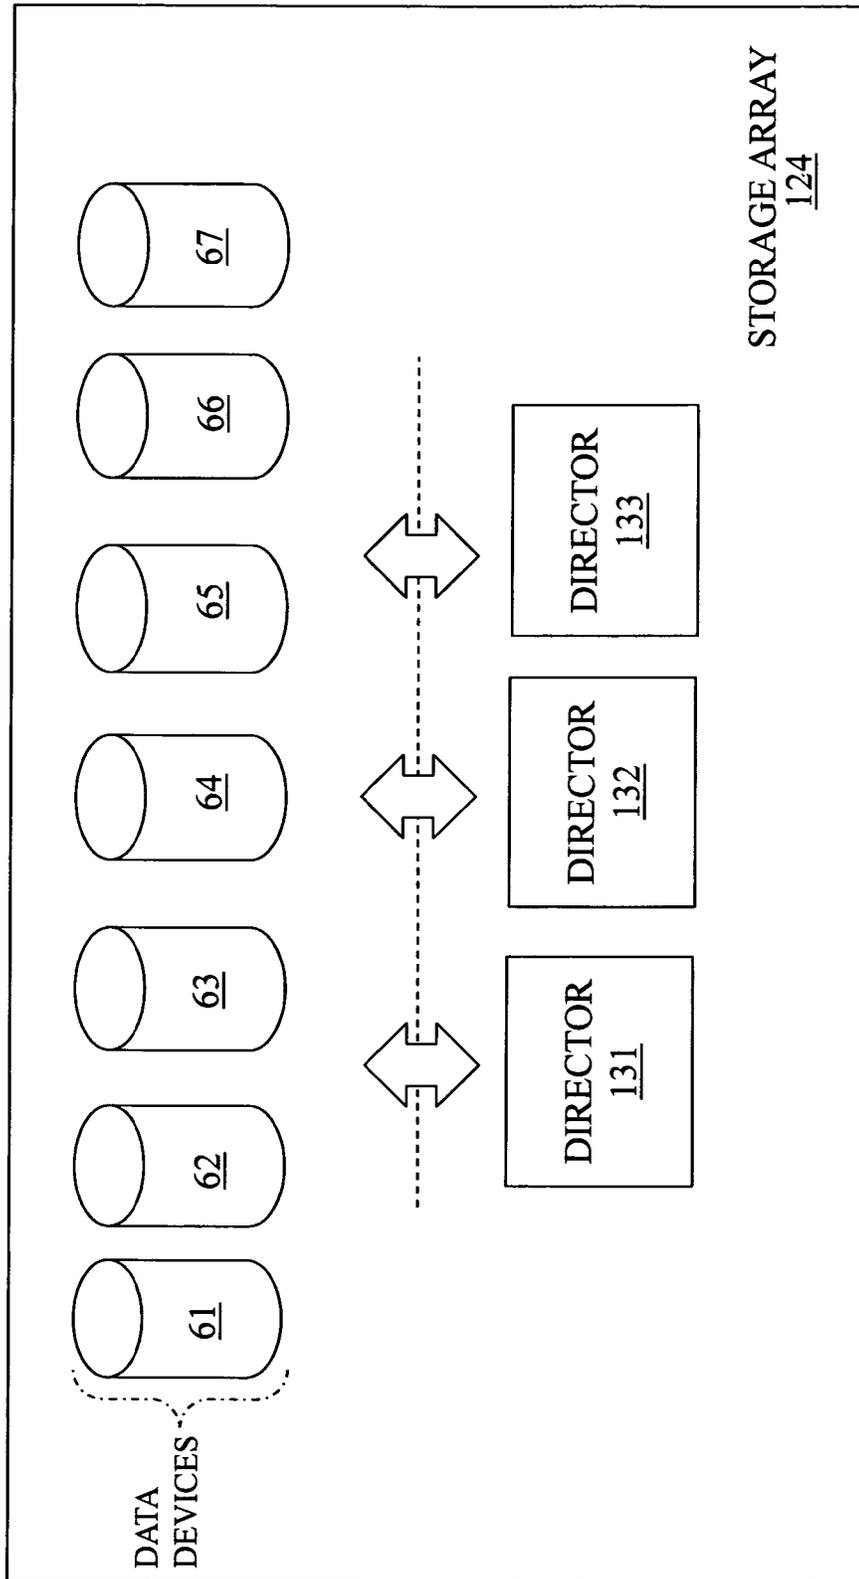


FIG. 5A

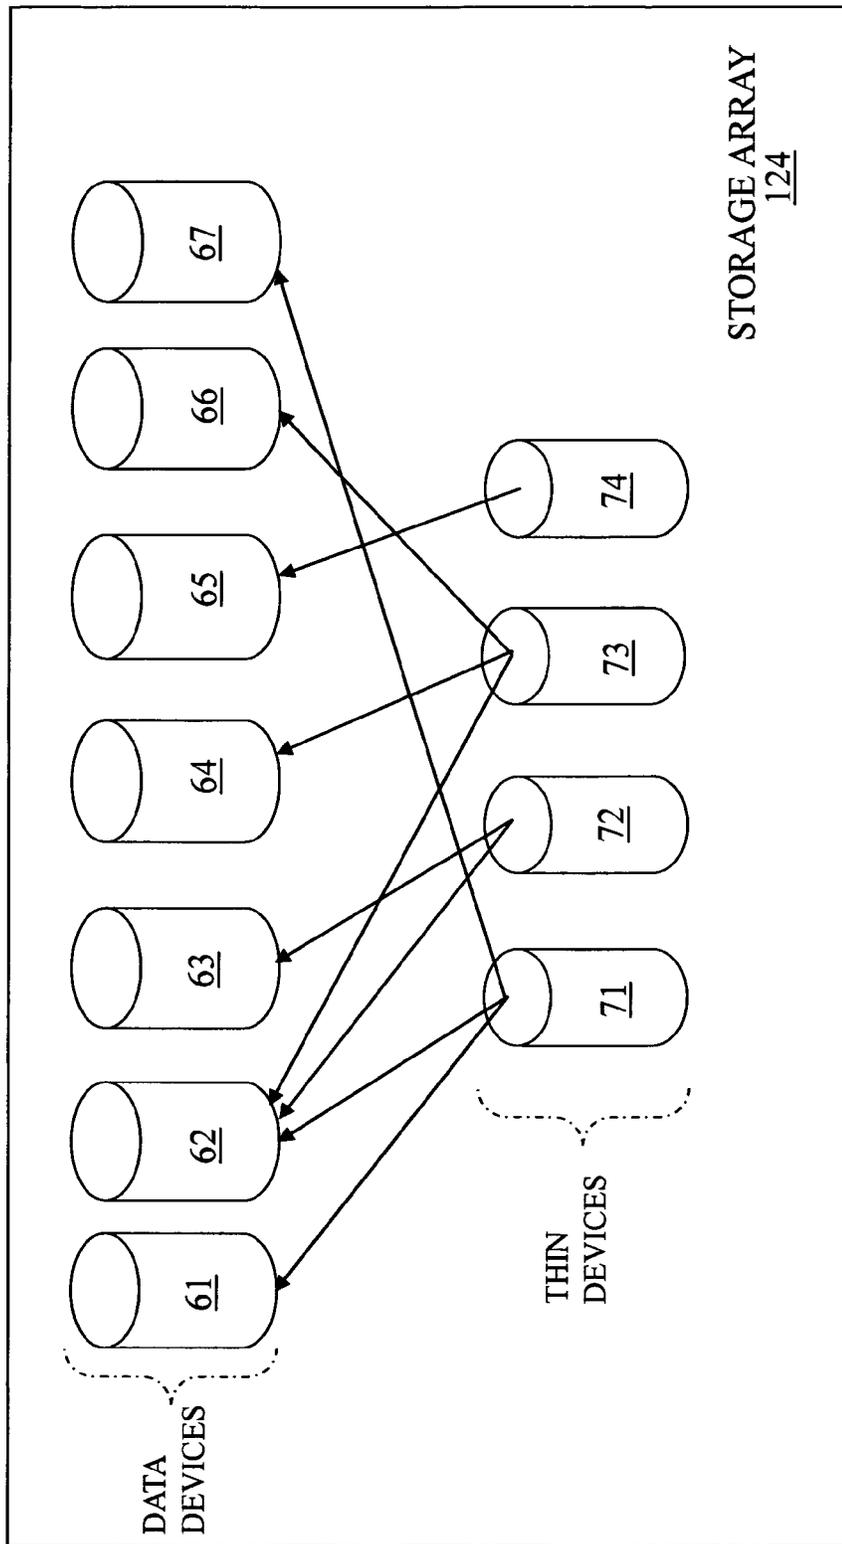


FIG. 5B

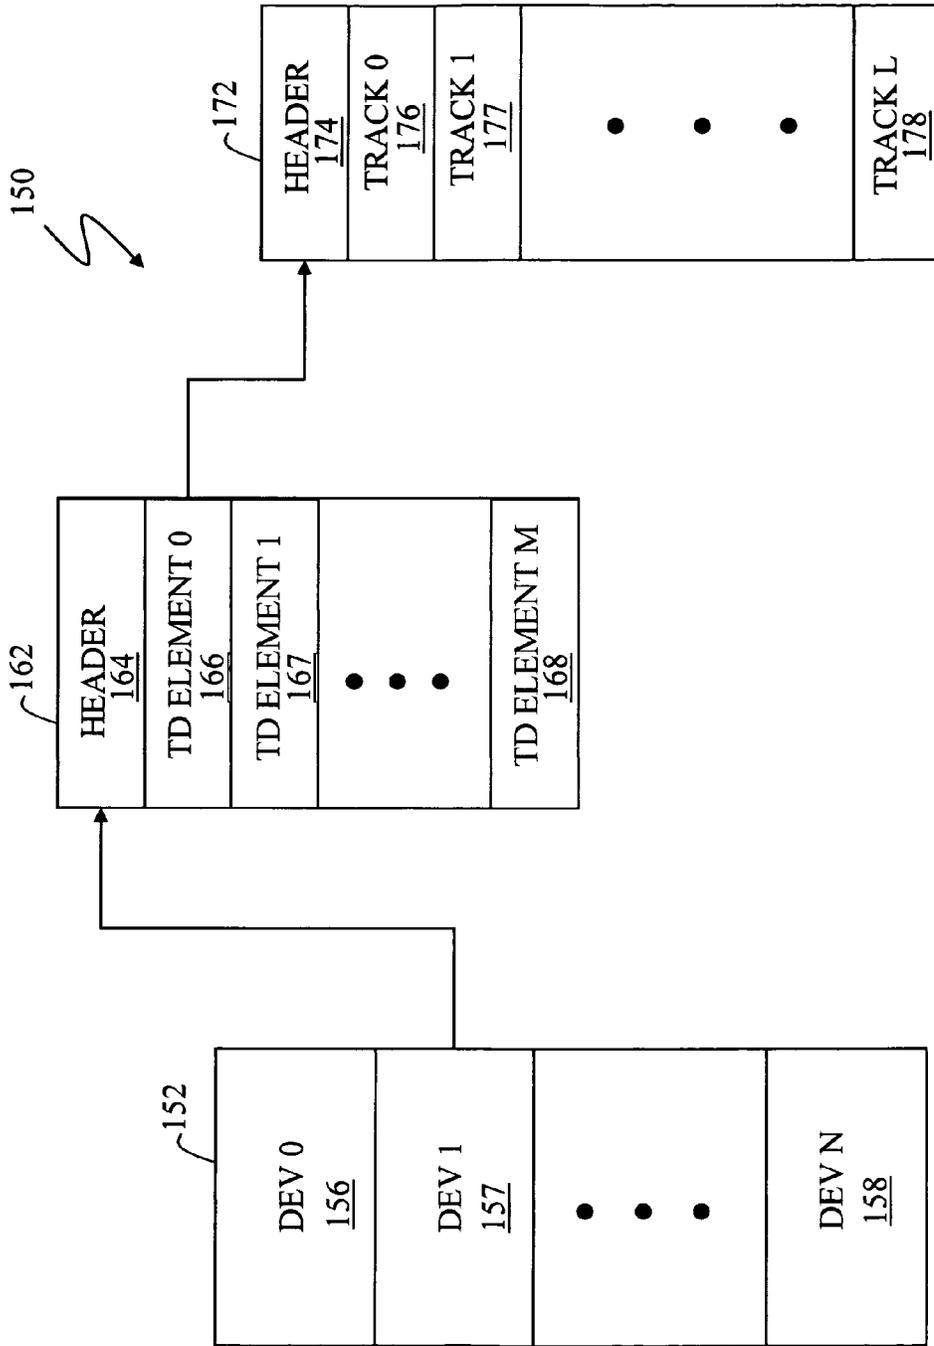


FIG 5C

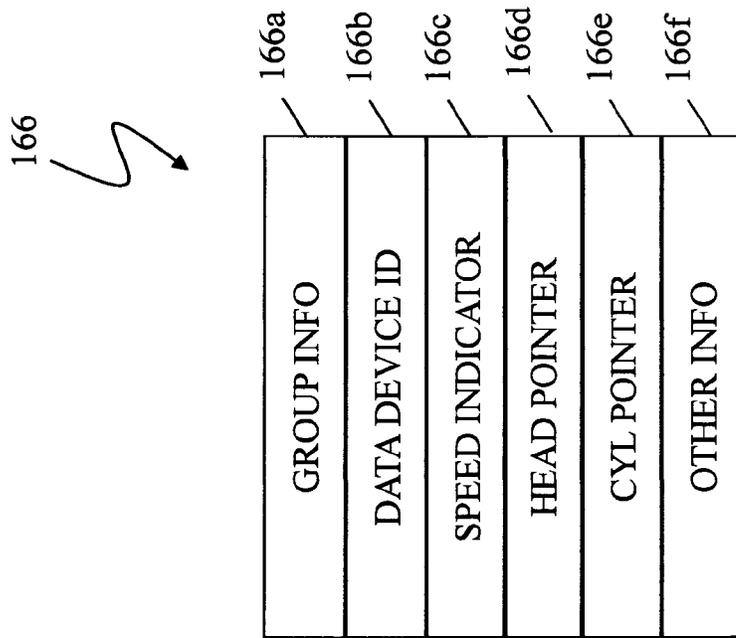


FIG 5D

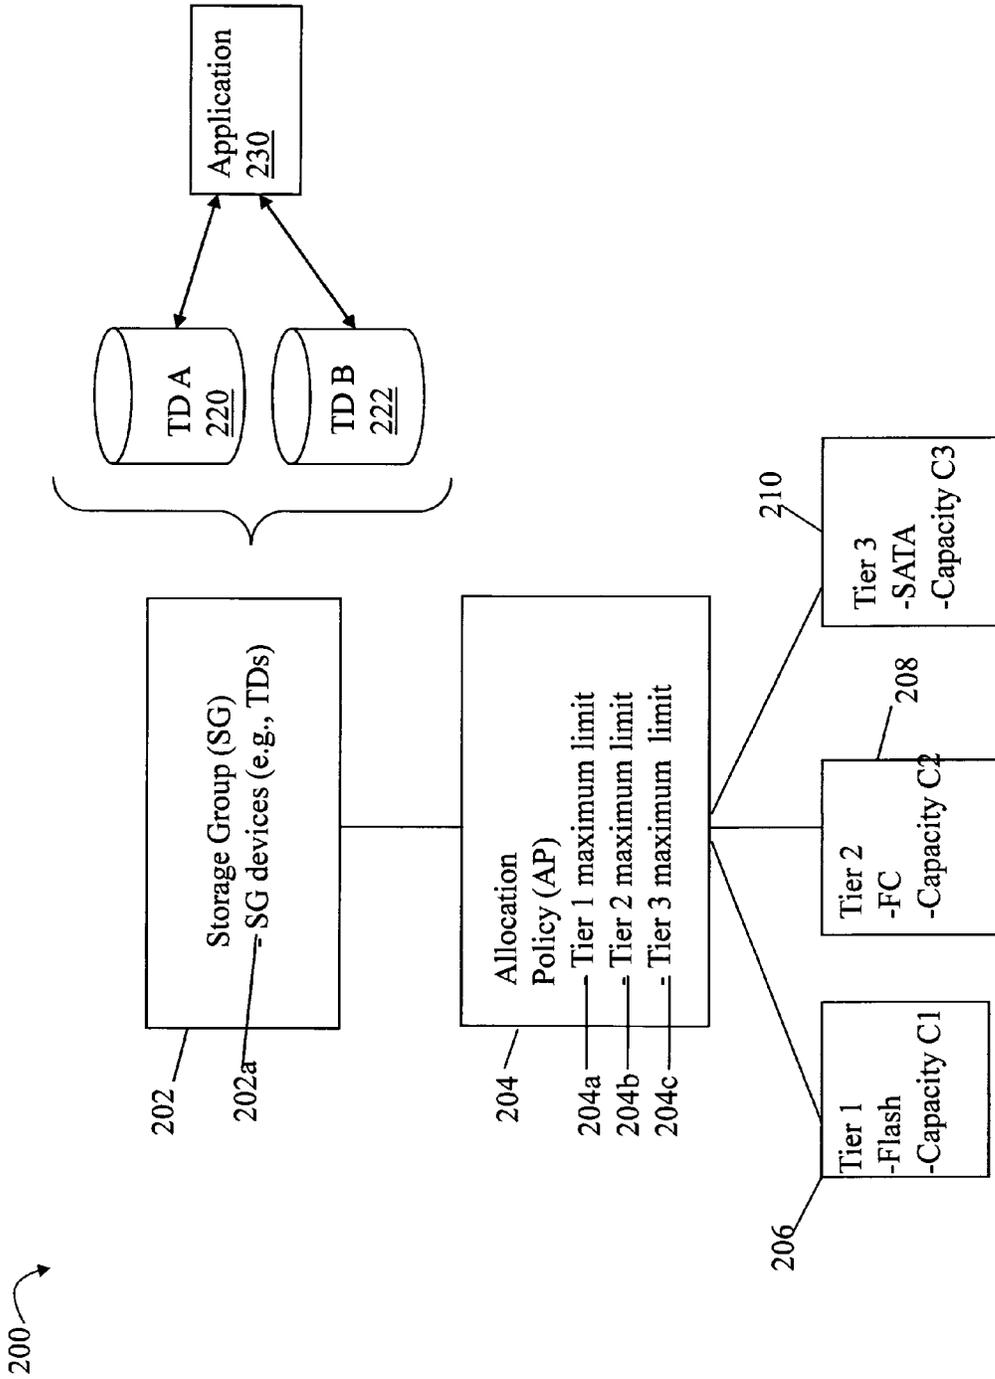


FIG. 6

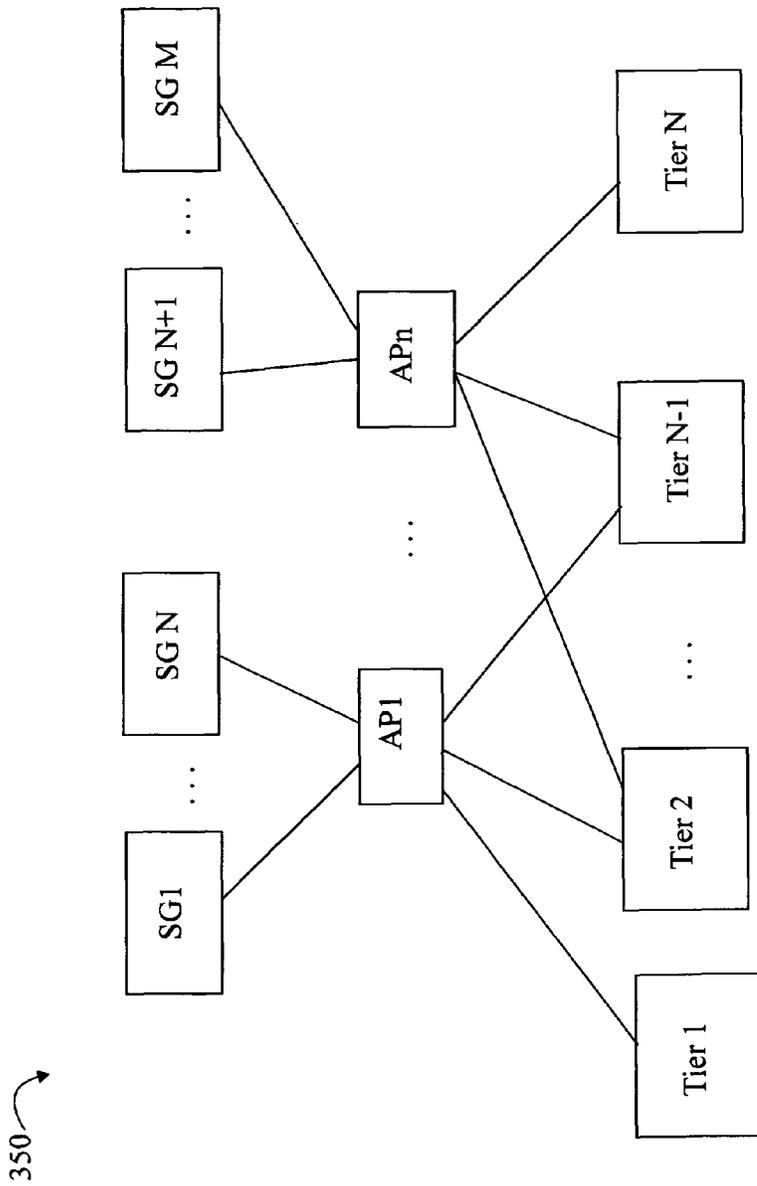
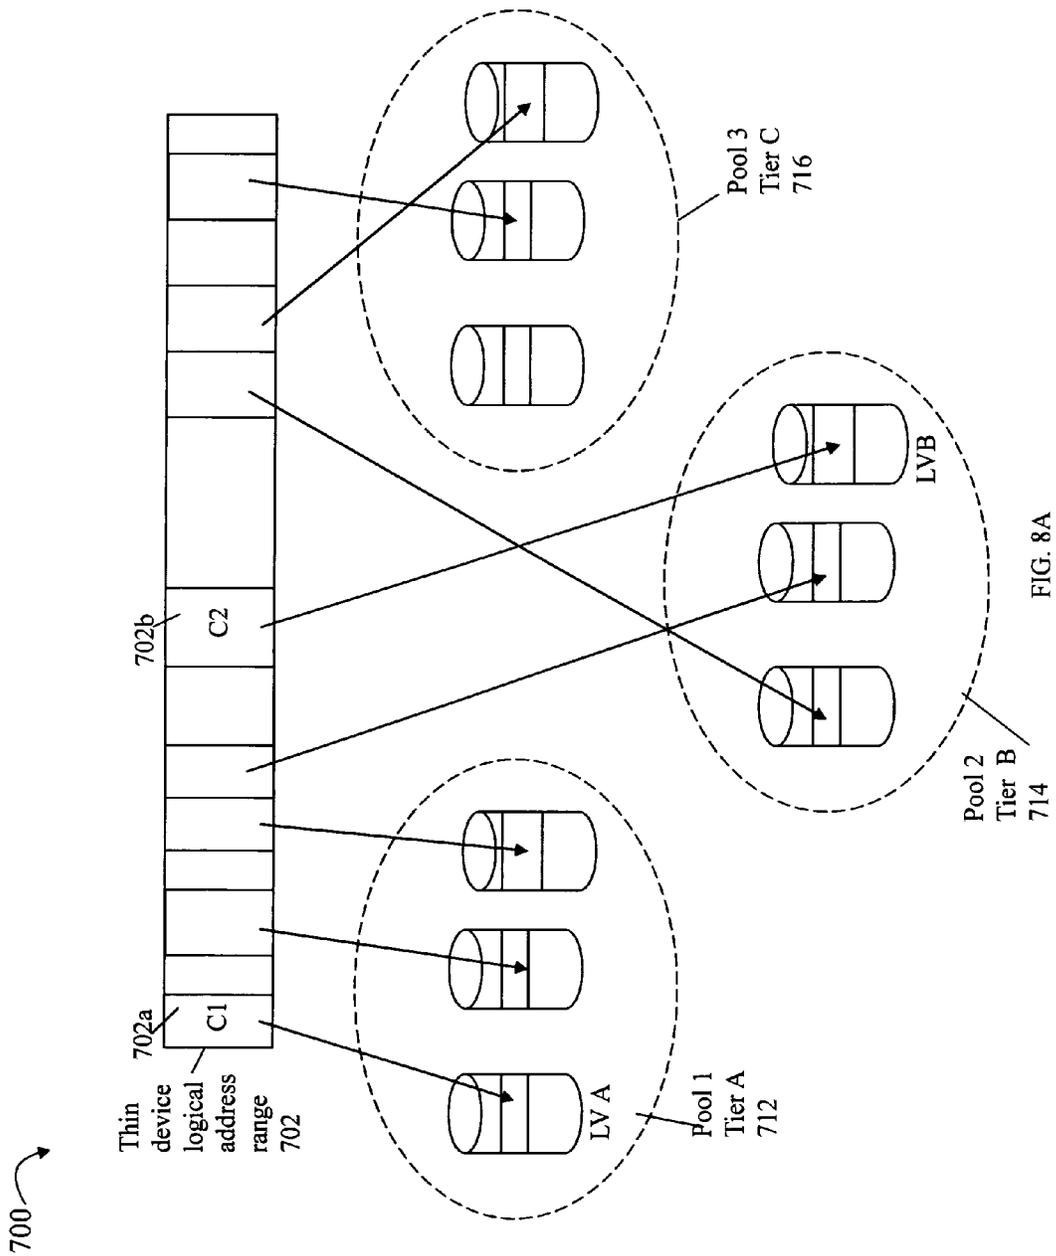


FIG. 7



750 →

Allocation map for TD1 760

760a Chunk ID	760b LV, starting offset	760c Storage Pool
C1	LV A, 100	Pool 1, Tier A
C2	LVB, 200	Pool 2, Tier B
:	:	:
:	:	:

762 →

764 →

FIG. 8B

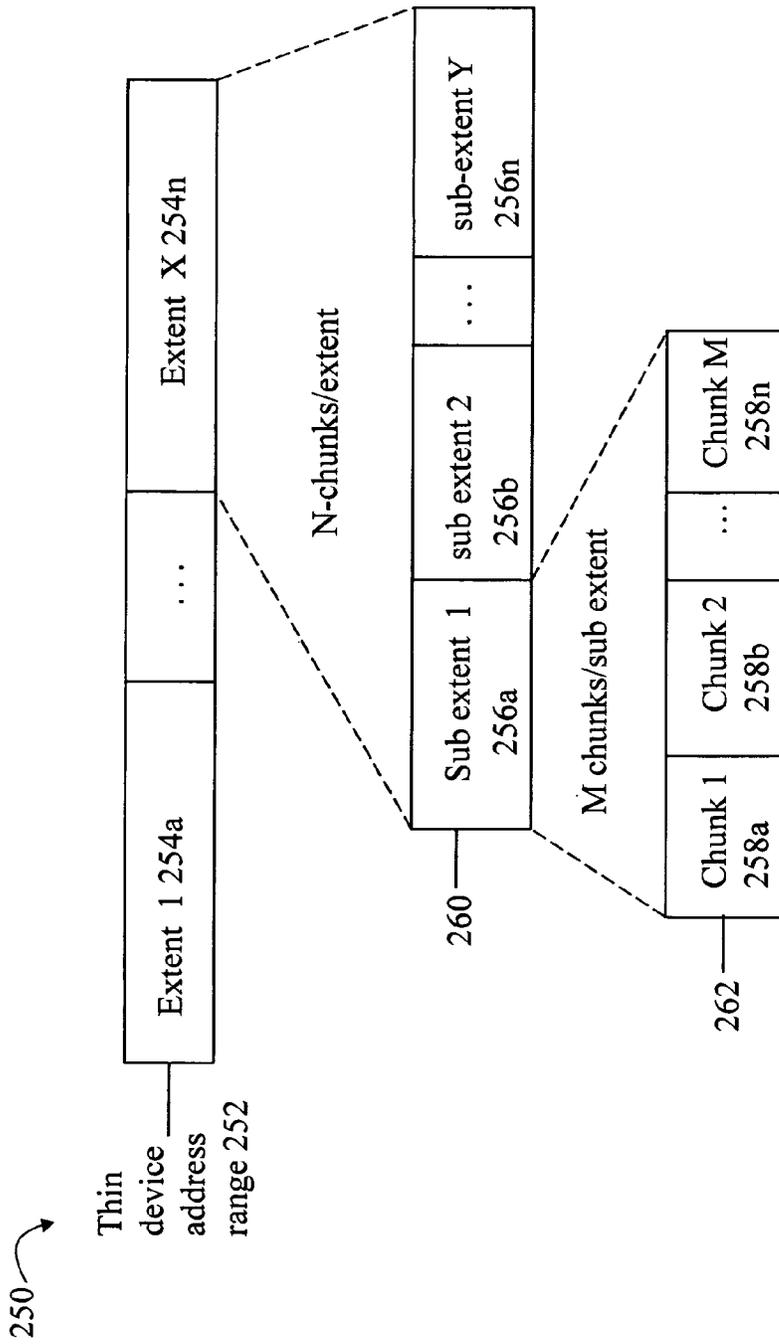


FIG. 9

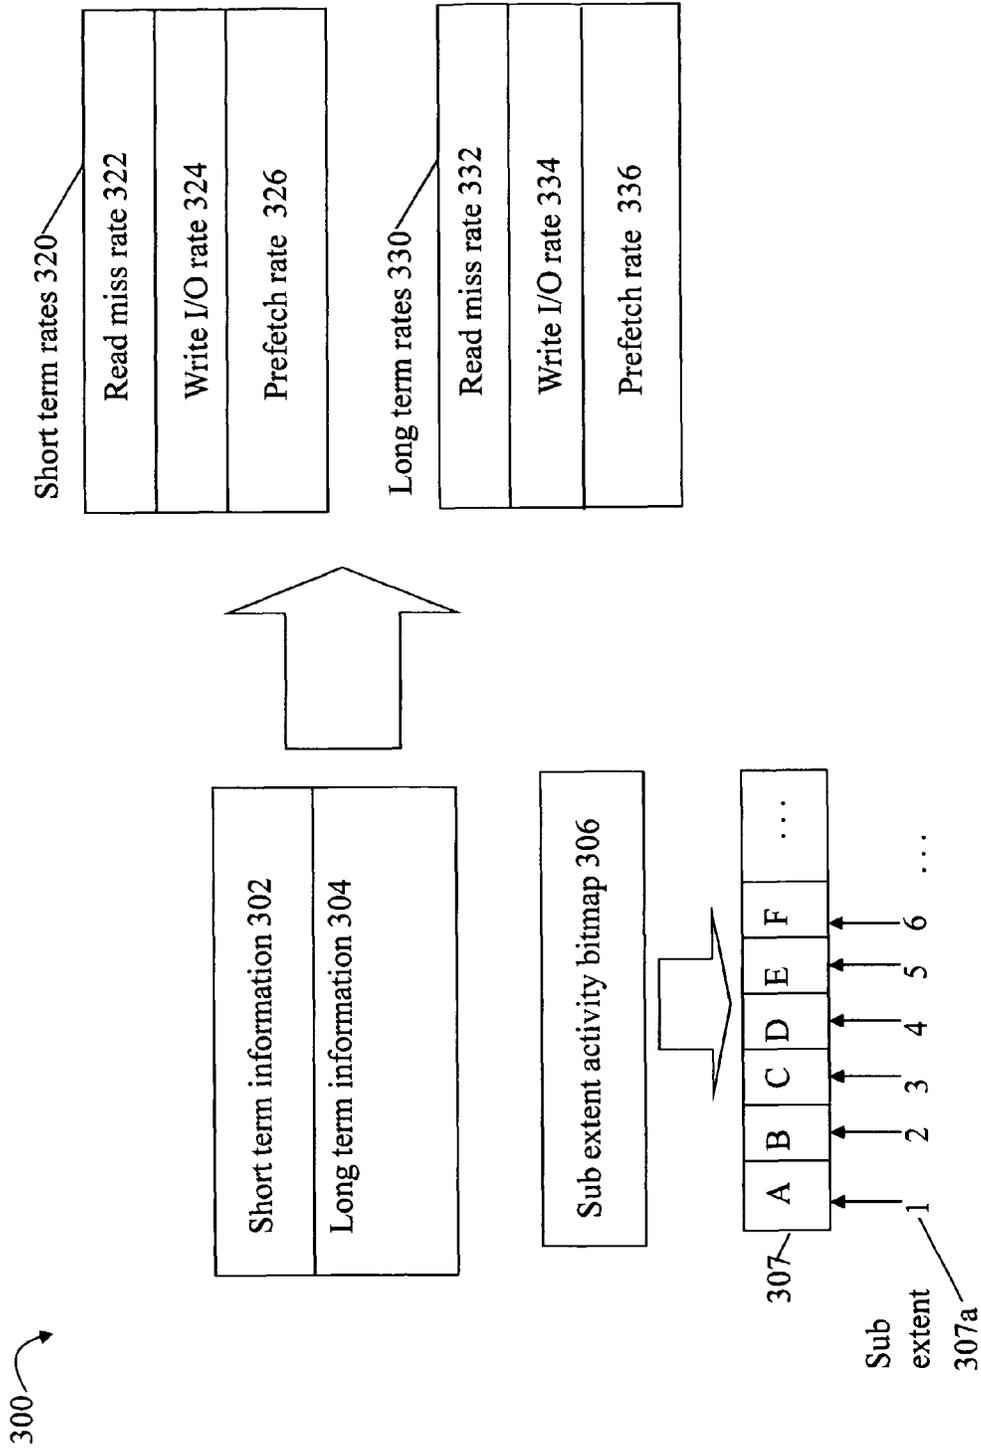


FIG. 10

400 ↗

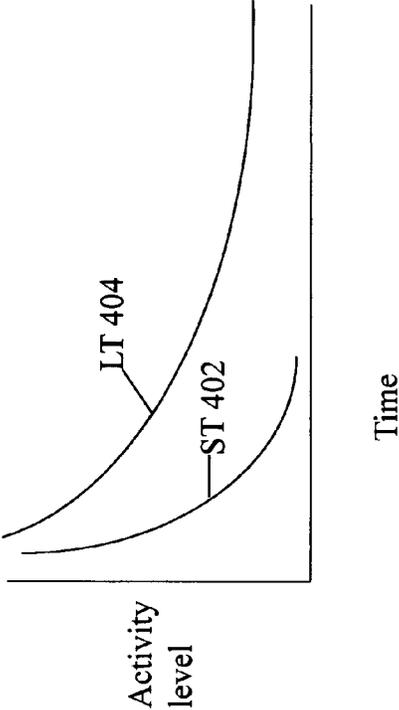


FIG. 11

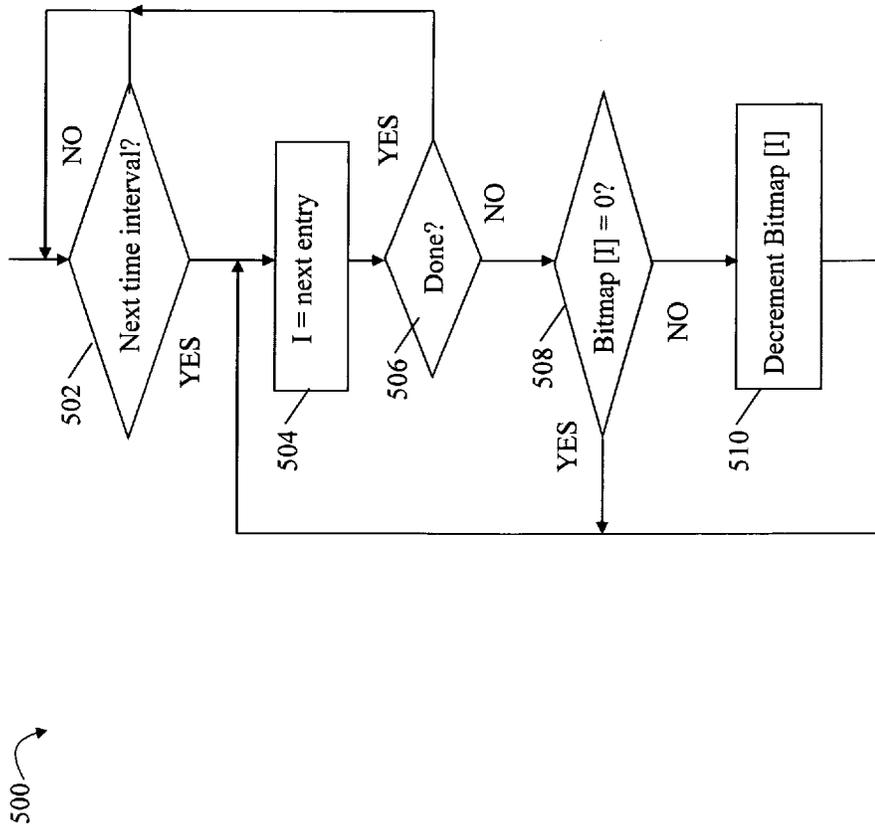


FIG. 12

550

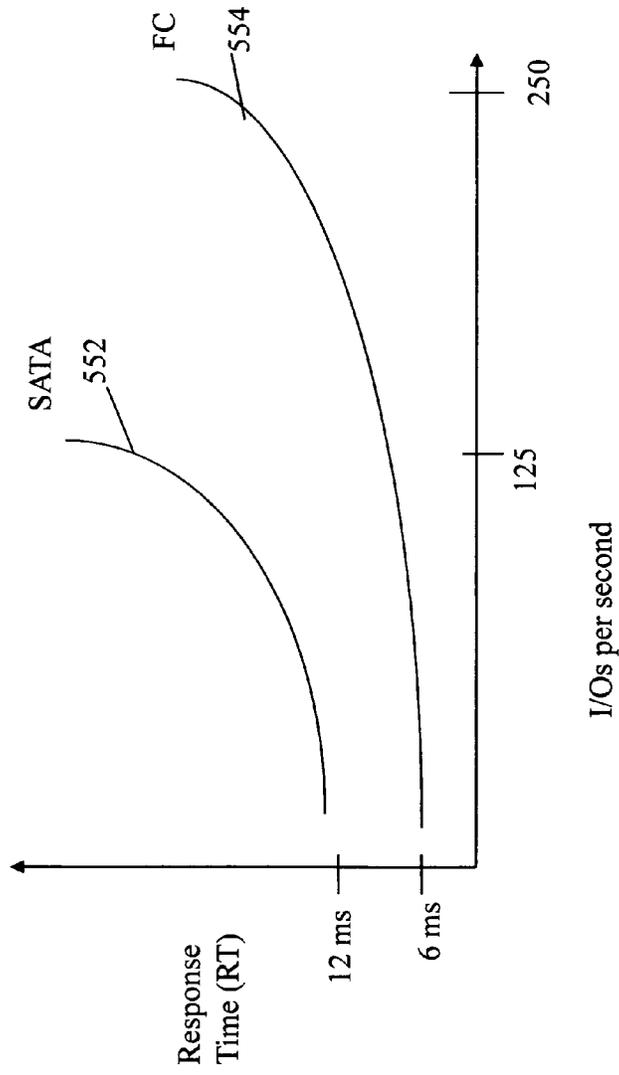


FIG. 13

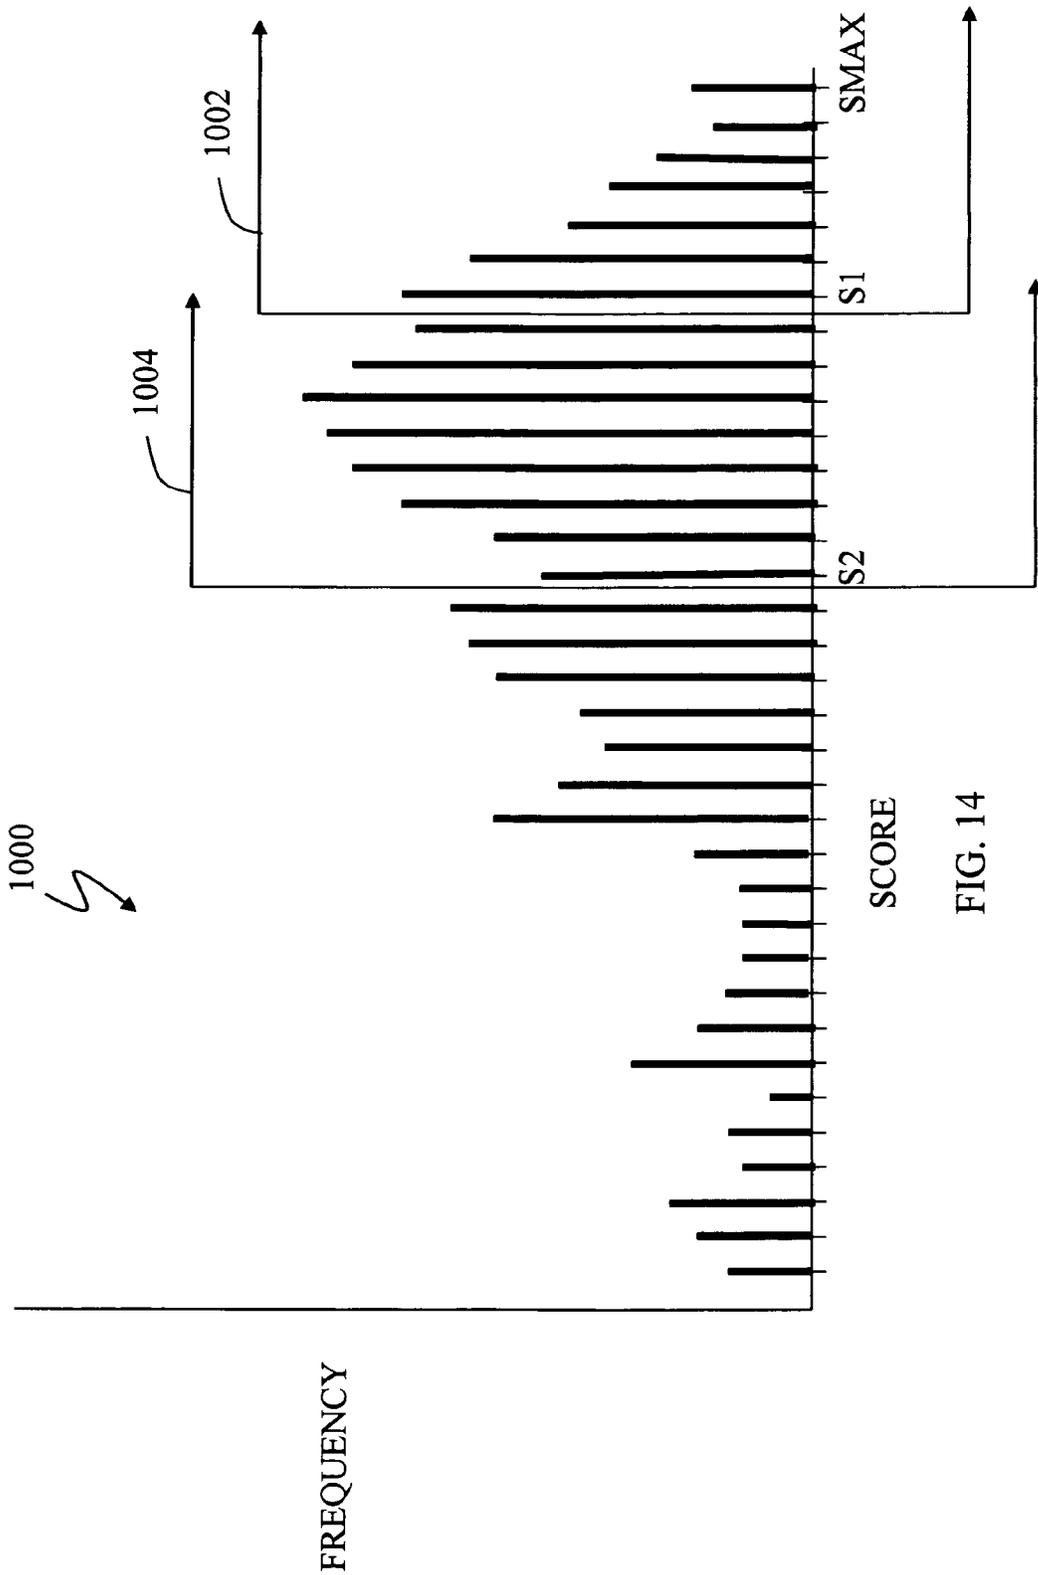


FIG. 14

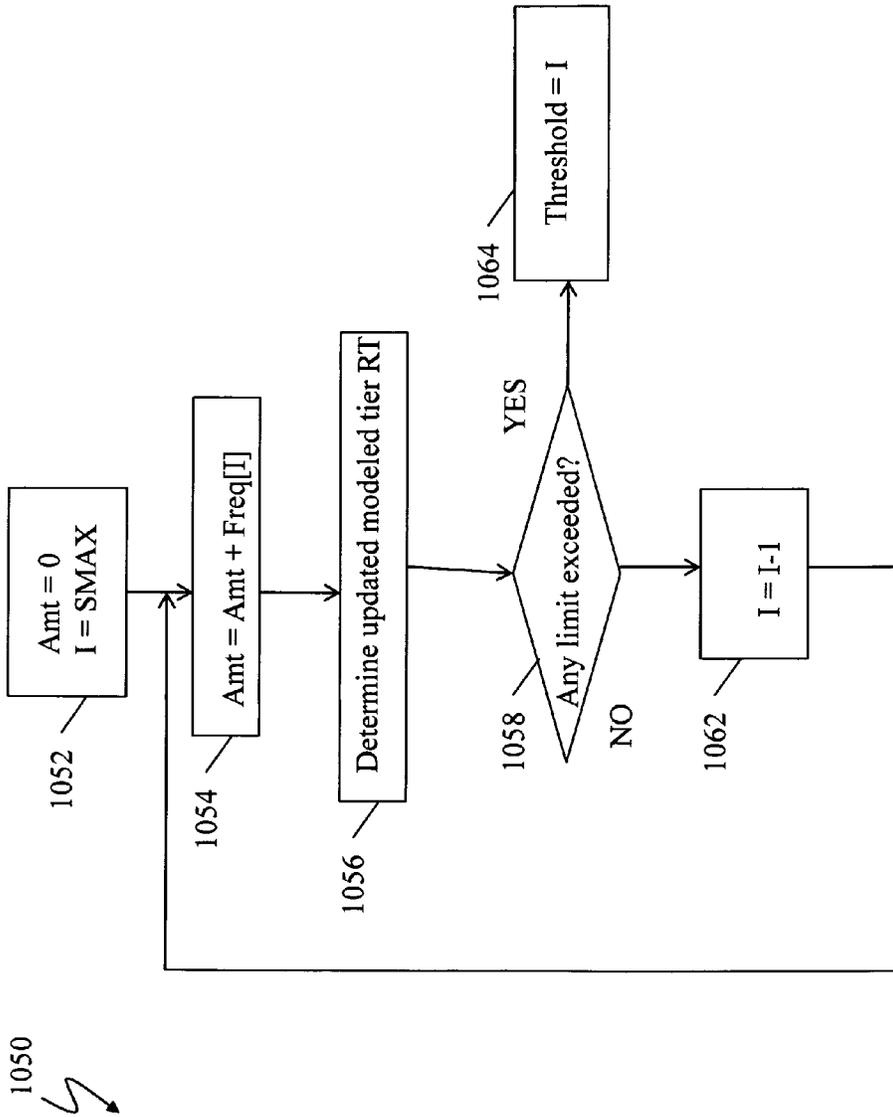


FIG. 15

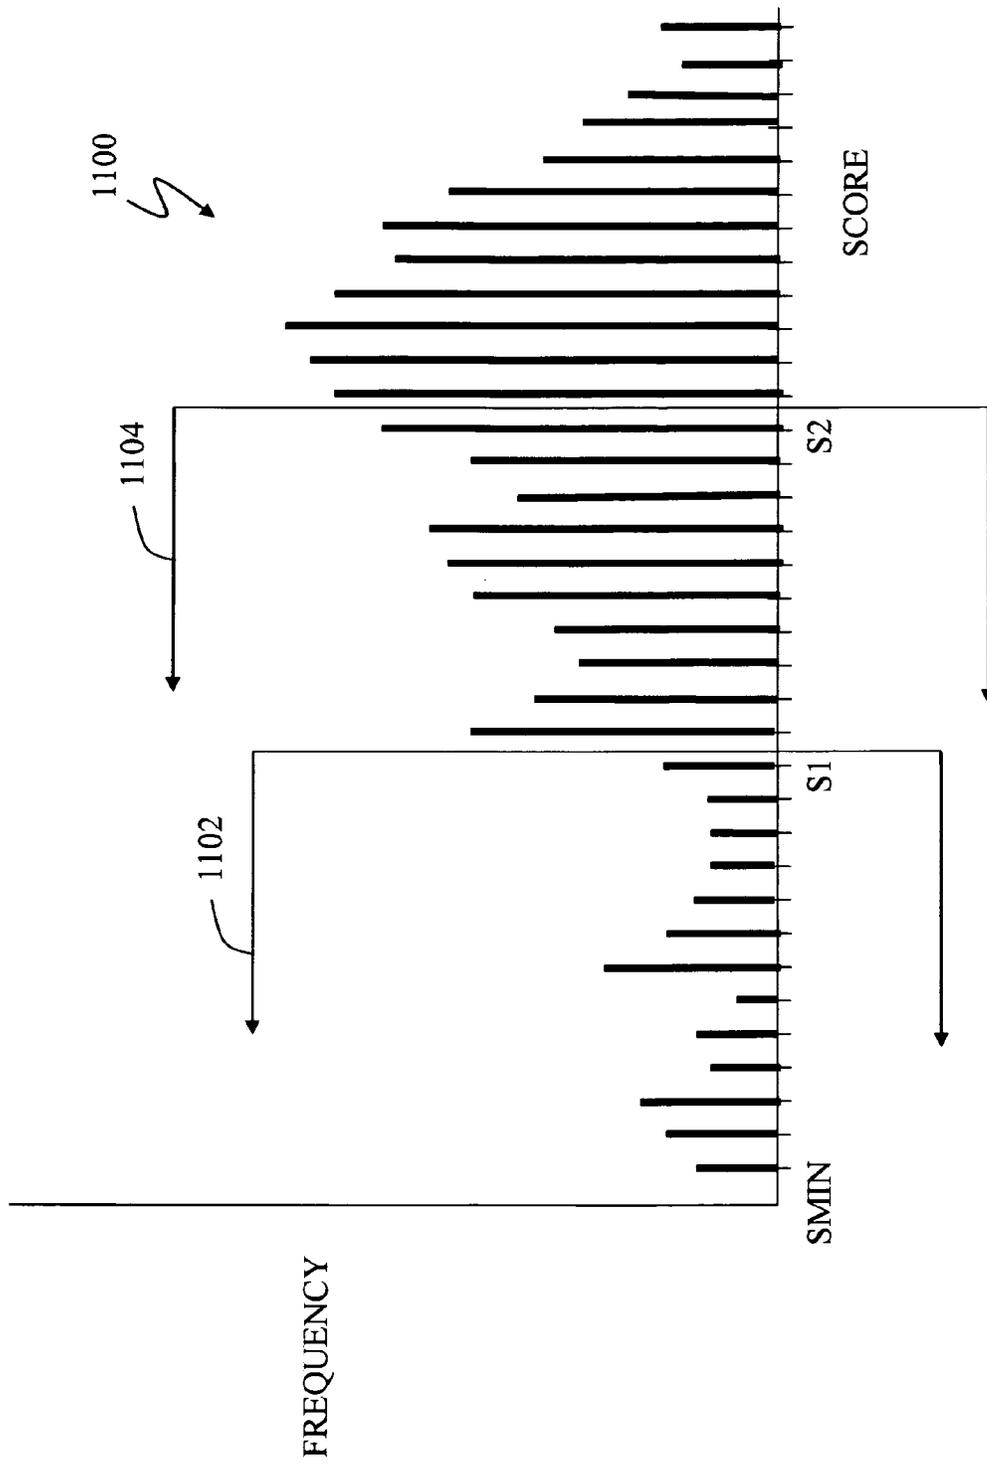


FIG. 16

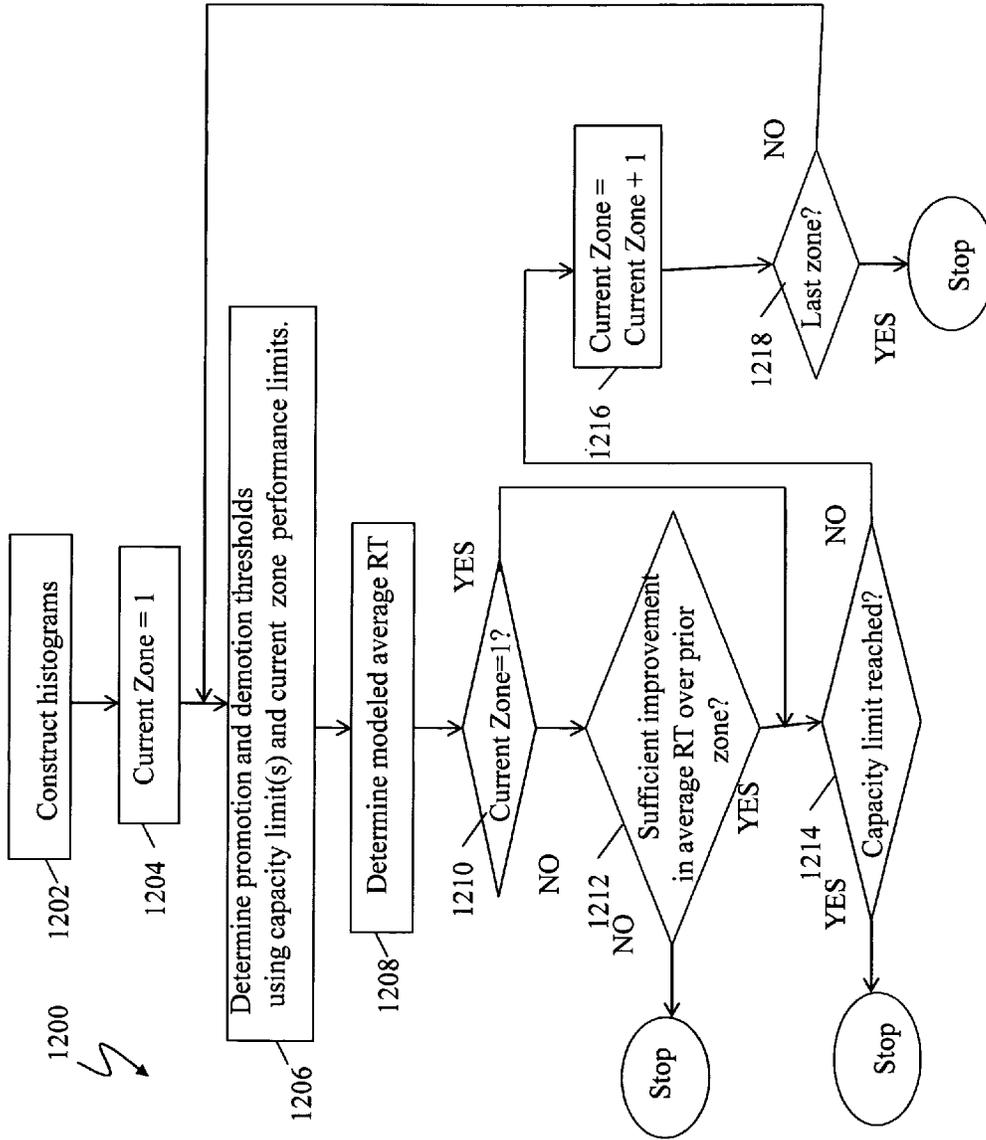


FIG. 17

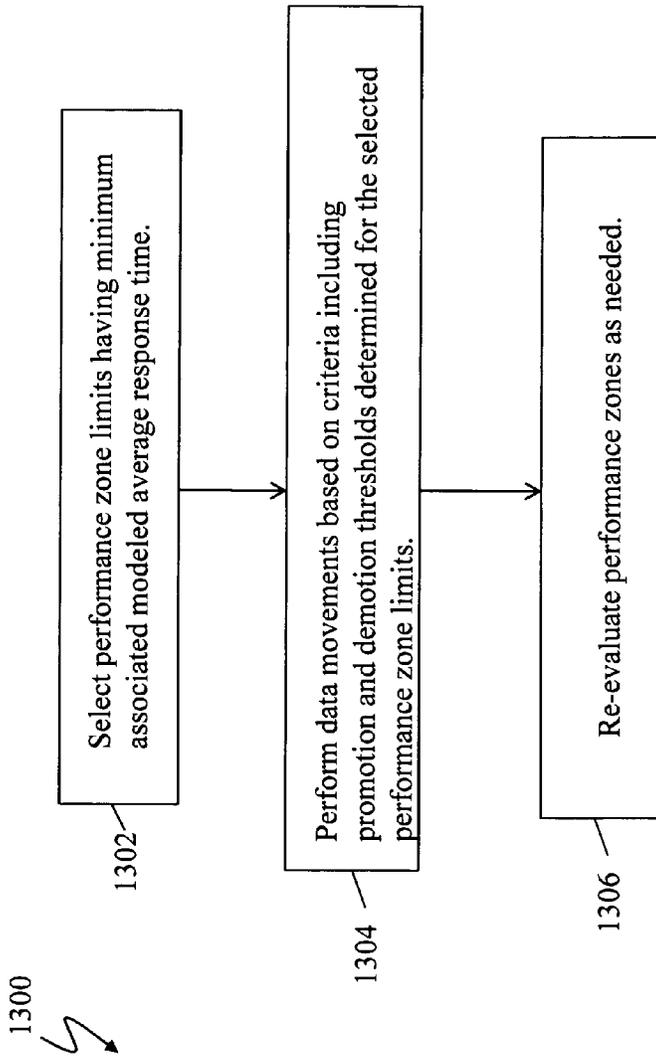


FIG. 18

**TECHNIQUES FOR AUTOMATED
EVALUATION AND MOVEMENT OF DATA
BETWEEN STORAGE TIERS FOR THIN
DEVICES**

BACKGROUND

1. Technical Field

This application generally relates to data storage, and more particularly to techniques used in connection with data storage configuration.

2. Description of Related Art

Computer systems may include different resources used by one or more host processors. Resources and host processors in a computer system may be interconnected by one or more communication connections. These resources may include, for example, data storage devices such as those included in the data storage systems manufactured by EMC Corporation. These data storage systems may be coupled to one or more host processors and provide storage services to each host processor. Multiple data storage systems from one or more different vendors may be connected and may provide common data storage for one or more host processors in a computer system.

A host processor may perform a variety of data processing tasks and operations using the data storage system. For example, a host processor may perform basic system I/O operations in connection with data requests, such as data read and write operations.

Host processor systems may store and retrieve data using a storage device containing a plurality of host interface units, disk drives, and disk interface units. Such storage devices are provided, for example, by EMC Corporation of Hopkinton, Mass. The host systems access the storage device through a plurality of channels provided therewith. Host systems provide data and access control information through the channels to the storage device and storage device provides data to the host systems also through the channels. The host systems do not address the disk drives of the storage device directly, but rather, access what appears to the host systems as a plurality of logical disk units, logical devices, or logical volumes (LVs). The logical disk units may or may not correspond to the actual disk drives. Allowing multiple host systems to access the single storage device unit allows the host systems to share data stored therein.

In connection with data storage, a variety of different technologies may be used. Data may be stored, for example, on different types of disk devices and/or flash memory devices. The data storage environment may define multiple storage tiers in which each tier includes physical devices or drives of varying technologies, performance characteristics, and the like. The physical devices of a data storage system, such as a data storage array, may be used to store data for multiple applications.

SUMMARY OF THE INVENTION

In accordance with one aspect of the invention is a method for evaluating data movement alternatives comprising receiving a set of criteria including capacity limits and performance limits, said performance limits including a plurality of sets of performance limits, each of said plurality of sets specifying a plurality of performance limits for a plurality of storage tiers; performing first processing to evaluate a plurality of alternatives for use in data movement with respect to a set of logical devices having data stored on a set of physical storage devices, each of said plurality of alternatives including a

different set of data movement criteria comprising said capacity limits and a different one of said plurality of sets of performance limits, said set of physical storage devices comprising at least a first physical device of one of said plurality of storage tiers and a second physical device of another one of said plurality of storage tiers; and selecting one of said plurality of sets of performance limits in accordance with said first processing. Each of the plurality of storage tiers may include storage devices having any of different technology and different performance characteristics than other storage devices included in others of said plurality of storage tiers. The plurality of storage tiers may include three storage tiers, a first storage tier being a highest performing storage tier including solid state storage devices, a second storage tier being a middle performing storage tier including disk drives, and a third storage tier being a lowest performing storage tier including disk drives. The capacity limits may include a capacity limit for each of the plurality of storage tiers. The capacity limits may include a capacity limit for a defined storage group of said set of logical devices. The first processing may include determining a promotion threshold for each of said plurality of storage tiers. The first processing may include determining a demotion threshold for each of said plurality of storage tiers. The set of logical devices may include one or more logical devices have data portions of the one or more logical devices stored on the set of physical storage devices. The logical devices may include one or more thin devices, each of said thin devices being a virtually provisioned device. Each of the thin devices may have a logical address range representing a presented storage capacity of said each thin device, and wherein at least a portion of said logical address range may not be mapped to physical storage indicating that physical storage is not allocated for said portion. The method may be performed in connection with optimization processing to optimize data storage system performance. The set of physical storage devices may be used to store data portions of said set of one or more logical devices and a second set of one or more logical devices. The set of logical devices may be defined to allow data movement between different ones of said physical storage devices and said second set of logical devices may be defined to not allow data movement between different ones of said physical storage devices. The method may include determining a performance baseline for said second set of logical devices, said performance baseline being defined as a measurement of workload of said second set of logical devices that is directed to said set of physical storage devices. Each of the plurality of sets of performance limits may include any one or more of a response time limit and a utilization limit for each of said plurality of storage tiers. The first processing may include performance modeling for each of said plurality of alternatives and determining an average response time for said set of physical devices. The first processing may include evaluating a first of said plurality of alternatives and proceeding to evaluate a second of said plurality of alternatives if said capacity limits have not been exceeded. Evaluation of the second alternatives may be performed if modeling of said first alternative indicates that said first alternative is determined to improve performance with respect to said set of physical devices by a threshold amount.

In accordance with another aspect of the invention is a computer readable medium comprising code stored thereon for evaluating data movement alternatives, the computer readable medium comprising code stored thereon for: receiving a set of criteria including capacity limits and performance limits, said performance limits including a plurality of sets of performance limits, each of said plurality of sets specifying a

3

plurality of performance limits for a plurality of storage tiers; performing first processing to evaluate a plurality of alternatives for use in data movement with respect to a set of logical devices having data stored on a set of physical storage devices, each of said plurality of alternatives including a different set of data movement criteria comprising said capacity limits and a different one of said plurality of sets of performance limits, said set of physical storage devices comprising at least a first physical device of one of said plurality of storage tiers and a second physical device of another one of said plurality of storage tiers; and selecting one of said plurality of sets of performance limits in accordance with said first processing. Each of the plurality of storage tiers may include storage devices having any of different technology and different performance characteristics than other storage devices included in others of said plurality of storage tiers. The plurality of storage tiers may include three storage tiers, a first storage tier being a highest performing storage tier including solid state storage devices, a second storage tier being a middle performing storage tier including disk drives, and a third storage tier being a lowest performing storage tier including disk drives.

BRIEF DESCRIPTION OF THE DRAWINGS

Features and advantages of the present invention will become more apparent from the following detailed description of exemplary embodiments thereof taken in conjunction with the accompanying drawings in which:

FIG. 1 is an example of an embodiment of a system that may utilize the techniques described herein;

FIG. 2 is a representation of the logical internal communications between the directors and memory included in one embodiment of a data storage system of FIG. 1;

FIG. 3 is an example representing components that may be included in a service processor in an embodiment in accordance with techniques herein;

FIGS. 4, 5A and 5B are examples illustrating a data storage system, such as data storage array, including a plurality of storage tiers in an embodiment in accordance with techniques herein;

FIG. 5C is a schematic diagram illustrating tables that are used to keep track of device information in connection with an embodiment of the system described herein;

FIG. 5D is a schematic diagram showing a group element of a thin device table in connection with an embodiment of the system described herein;

FIGS. 6 and 7 are examples illustrating a storage group, allocation policy and associated storage tiers in an embodiment in accordance with techniques herein;

FIGS. 8A and 8B are examples illustrating thin devices and associated structures that may be used in an embodiment in accordance with techniques herein;

FIG. 9 is an example illustrating data portions comprising a thin device's logical address range;

FIG. 10 is an example of performance information that may be determined in connection with thin devices in an embodiment in accordance with techniques herein;

FIG. 11 is a graphical illustration of long term and short term statistics described herein;

FIGS. 12, 15, 17 and 18 are flowcharts of processing steps that may be performed in an embodiment in accordance with techniques herein;

4

FIG. 13 is an example of performance curves used to model device response time in an embodiment in accordance with techniques herein; and

FIGS. 14 and 16 illustrate histograms that may be used in threshold selection in accordance with techniques herein.

DETAILED DESCRIPTION OF EMBODIMENT(S)

Referring to FIG. 1, shown is an example of an embodiment of a system that may be used in connection with performing the techniques described herein. The system 10 includes a data storage system 12 connected to host systems 14a-14n through communication medium 18. In this embodiment of the computer system 10, and the n hosts 14a-14n may access the data storage system 12, for example, in performing input/output (I/O) operations or data requests. The communication medium 18 may be any one or more of a variety of networks or other type of communication connections as known to those skilled in the art. The communication medium 18 may be a network connection, bus, and/or other type of data link, such as a hardwire or other connections known in the art. For example, the communication medium 18 may be the Internet, an intranet, network (including a Storage Area Network (SAN)) or other wireless or other hardwired connection(s) by which the host systems 14a-14n may access and communicate with the data storage system 12, and may also communicate with other components included in the system 10.

Each of the host systems 14a-14n and the data storage system 12 included in the system 10 may be connected to the communication medium 18 by any one of a variety of connections as may be provided and supported in accordance with the type of communication medium 18. The processors included in the host computer systems 14a-14n may be any one of a variety of proprietary or commercially available single or multi-processor system, such as an Intel-based processor, or other type of commercially available processor able to support traffic in accordance with each particular embodiment and application.

It should be noted that the particular examples of the hardware and software that may be included in the data storage system 12 are described herein in more detail, and may vary with each particular embodiment. Each of the host computers 14a-14n and data storage system may all be located at the same physical site, or, alternatively, may also be located in different physical locations. Examples of the communication medium that may be used to provide the different types of connections between the host computer systems and the data storage system of the system 10 may use a variety of different communication protocols such as SCSI, Fibre Channel, iSCSI, and the like. Some or all of the connections by which the hosts and data storage system may be connected to the communication medium may pass through other communication devices, such as a Connectrix or other switching equipment that may exist such as a phone line, a repeater, a multiplexer or even a satellite.

Each of the host computer systems may perform different types of data operations in accordance with different types of tasks. In the embodiment of FIG. 1, any one of the host computers 14a-14n may issue a data request to the data storage system 12 to perform a data operation. For example, an application executing on one of the host computers 14a-14n may perform a read or write operation resulting in one or more data requests to the data storage system 12.

It should be noted that although element 12 is illustrated as a single data storage system, such as a single data storage

array, element **12** may also represent, for example, multiple data storage arrays alone, or in combination with, other data storage devices, systems, appliances, and/or components having suitable connectivity, such as in a SAN, in an embodiment using the techniques herein. It should also be noted that an embodiment may include data storage arrays or other components from one or more vendors. In subsequent examples illustrated the techniques herein, reference may be made to a single data storage array by a vendor, such as by EMC Corporation of Hopkinton, Mass. However, as will be appreciated by those skilled in the art, the techniques herein are applicable for use with other data storage arrays by other vendors and with other components than as described herein for purposes of example.

The data storage system **12** may be a data storage array including a plurality of data storage devices **16a-16n**. The data storage devices **16a-16n** may include one or more types of data storage devices such as, for example, one or more disk drives and/or one or more solid state drives (SSDs). An SSD is a data storage device that uses solid-state memory to store persistent data. An SSD using SRAM or DRAM, rather than flash memory, may also be referred to as a RAM drive. SSD may refer to solid state electronics devices as distinguished from electromechanical devices, such as hard drives, having moving parts. Flash devices or flash memory-based SSDs are one type of SSD that contains no moving parts. As described in more detail in following paragraphs, the techniques herein may be used in an embodiment in which one or more of the devices **16a-16n** are flash drives or devices. More generally, the techniques herein may also be used with any type of SSD although following paragraphs may make reference to a particular type such as a flash device or flash memory device.

The data storage array may also include different types of adapters or directors, such as an HA **21** (host adapter), RA **40** (remote adapter), and/or device interface **23**. Each of the adapters may be implemented using hardware including a processor with local memory with code stored thereon for execution in connection with performing different operations. The HAs may be used to manage communications and data operations between one or more host systems and the global memory (GM). In an embodiment, the HA may be a Fibre Channel Adapter (FA) or other adapter which facilitates host communication. The HA **21** may be characterized as a front end component of the data storage system which receives a request from the host. The data storage array may include one or more RAs that may be used, for example, to facilitate communications between data storage arrays. The data storage array may also include one or more device interfaces **23** for facilitating data transfers to/from the data storage devices **16a-16n**. The data storage interfaces **23** may include device interface modules, for example, one or more disk adapters (DAs) (e.g., disk controllers), adapters used to interface with the flash drives, and the like. The DAs may also be characterized as back end components of the data storage system which interface with the physical data storage devices.

One or more internal logical communication paths may exist between the device interfaces **23**, the RAs **40**, the HAs **21**, and the memory **26**. An embodiment, for example, may use one or more internal busses and/or communication modules. For example, the global memory portion **25b** may be used to facilitate data transfers and other communications between the device interfaces, HAs and/or RAs in a data storage array. In one embodiment, the device interfaces **23** may perform data operations using a cache that may be included in the global memory **25b**, for example, when communicating with other device interfaces and other compo-

nents of the data storage array. The other portion **25a** is that portion of memory that may be used in connection with other designations that may vary in accordance with each embodiment.

The particular data storage system as described in this embodiment, or a particular device thereof, such as a disk or particular aspects of a flash device, should not be construed as a limitation. Other types of commercially available data storage systems, as well as processors and hardware controlling access to these particular devices, may also be included in an embodiment.

Host systems provide data and access control information through channels to the storage systems, and the storage systems may also provide data to the host systems also through the channels. The host systems do not address the drives or devices **16a-16n** of the storage systems directly, but rather access to data may be provided to one or more host systems from what the host systems view as a plurality of logical devices or logical volumes (LVs). The LVs may or may not correspond to the actual physical devices or drives **16a-16n**. For example, one or more LVs may reside on a single physical drive or multiple drives. Data in a single data storage system, such as a single data storage array, may be accessed by multiple hosts allowing the hosts to share the data residing therein. The HAs may be used in connection with communications between a data storage array and a host system. The RAs may be used in facilitating communications between two data storage arrays. The DAs may be one type of device interface used in connection with facilitating data transfers to/from the associated disk drive(s) and LV(s) residing thereon. A flash device interface may be another type of device interface used in connection with facilitating data transfers to/from the associated flash devices and LV(s) residing thereon. It should be noted that an embodiment may use the same or a different device interface for one or more different types of devices than as described herein.

The device interface, such as a DA, performs I/O operations on a drive **16a-16n**. In the following description, data residing on an LV may be accessed by the device interface following a data request in connection with I/O operations that other directors originate. Data may be accessed by LV in which a single device interface manages data requests in connection with the different one or more LVs that may reside on a drive **16a-16n**. For example, a device interface may be a DA that accomplishes the foregoing by creating job records for the different LVs associated with a particular device. These different job records may be associated with the different LVs in a data structure stored and managed by each device interface.

Also shown in FIG. 1 is a service processor **22a** that may be used to manage and monitor the system **12**. In one embodiment, the service processor **22a** may be used in collecting performance data, for example, regarding the I/O performance in connection with data storage system **12**. This performance data may relate to, for example, performance measurements in connection with a data request as may be made from the different host computer systems **14a 14n**. This performance data may be gathered and stored in a storage area. Additional detail regarding the service processor **22a** is described in following paragraphs.

It should be noted that a service processor **22a** may exist external to the data storage system **12** and may communicate with the data storage system **12** using any one of a variety of communication connections. In one embodiment, the service processor **22a** may communicate with the data storage system **12** through three different connections, a serial port, a parallel port and using a network interface card, for example, with an

Ethernet connection. Using the Ethernet connection, for example, a service processor may communicate directly with DAs and HAs within the data storage system **12**.

Referring to FIG. 2, shown is a representation of the logical internal communications between the directors and memory included in a data storage system. Included in FIG. 2 is a plurality of directors **37a-37n** coupled to the memory **26**. Each of the directors **37a-37n** represents one of the HAs, RAs, or device interfaces that may be included in a data storage system. In an embodiment disclosed herein, there may be up to sixteen directors coupled to the memory **26**. Other embodiments may allow a maximum number of directors other than sixteen as just described and the maximum number may vary with embodiment.

The representation of FIG. 2 also includes an optional communication module (CM) **38** that provides an alternative communication path between the directors **37a-37n**. Each of the directors **37a-37n** may be coupled to the CM **38** so that any one of the directors **37a-37n** may send a message and/or data to any other one of the directors **37a-37n** without needing to go through the memory **26**. The CM **38** may be implemented using conventional MUX/router technology where a sending one of the directors **37a-37n** provides an appropriate address to cause a message and/or data to be received by an intended receiving one of the directors **37a-37n**. In addition, a sending one of the directors **37a-37n** may be able to broadcast a message to all of the other directors **37a-37n** at the same time.

With reference back to FIG. 1, components of the data storage system may communicate using GM **25b**. For example, in connection with a write operation, an embodiment may first store the data in cache included in a portion of GM **25b**, mark the cache slot including the write operation data as write pending (WP), and then later destage the WP data from cache to one of the devices **16a-16n**. In connection with returning data to a host from one of the devices as part of a read operation, the data may be copied from the device by the appropriate device interface, such as a DA servicing the device. The device interface may copy the data read into a cache slot included in GM which is, in turn, communicated to the appropriate HA in communication with the host.

As described above, the data storage system **12** may be a data storage array including a plurality of data storage devices **16a-16n** in which one or more of the devices **16a-16n** are flash memory devices employing one or more different flash memory technologies. In one embodiment, the data storage system **12** may be a Symmetrix® DMX™ or VMAX™ data storage array by EMC Corporation of Hopkinton, Mass. In the foregoing data storage array, the data storage devices **16a-16n** may include a combination of disk devices and flash devices in which the flash devices may appear as standard Fibre Channel (FC) drives to the various software tools used in connection with the data storage array. The flash devices may be constructed using nonvolatile semiconductor NAND flash memory. The flash devices may include one or more SLC (single level cell) devices and/or MLC (multi level cell) devices.

It should be noted that the techniques herein may be used in connection with flash devices comprising what may be characterized as enterprise-grade or enterprise-class flash drives (EFDs) with an expected lifetime (e.g., as measured in an amount of actual elapsed time such as a number of years, months, and/or days) based on a number of guaranteed write cycles, or program cycles, and a rate or frequency at which the writes are performed. Thus, a flash device may be expected to have a usage measured in calendar or wall clock elapsed time based on the amount of time it takes to perform the number of

guaranteed write cycles. The techniques herein may also be used with other flash devices, more generally referred to as non-enterprise class flash devices, which, when performing writes at a same rate as for enterprise class drives, may have a lower expected lifetime based on a lower number of guaranteed write cycles.

The techniques herein may be generally used in connection with any type of flash device, or more generally, any SSD technology. The flash device may be, for example, a flash device which is a NAND gate flash device, NOR gate flash device, flash device that uses SLC or MLC technology, and the like, as known in the art. In one embodiment, the one or more flash devices may include MLC flash memory devices although an embodiment may utilize MLC, alone or in combination with, other types of flash memory devices or other suitable memory and data performance technologies. More generally, the techniques herein may be used in connection with other SSD technologies although particular flash memory technologies may be described herein for purposes of illustration.

An embodiment in accordance with techniques herein may have one or more defined storage tiers. Each tier may generally include physical storage devices or drives having one or more attributes associated with a definition for that tier. For example, one embodiment may provide a tier definition based on a set of one or more attributes. The attributes may include any one or more of a storage type or storage technology, a type of data protection, device performance characteristic(s), storage capacity, and the like. The storage type or technology may specify whether a physical storage device is an SSD drive (such as a flash drive), a particular type of SSD drive (such using flash or a form of RAM), a type of magnetic disk or other non-SSD drive (such as an FC disk drive, a SATA (Serial Advanced Technology Attachment) drive), and the like. Data protection may specify a type or level of data storage protection such, for example, as a particular RAID level (e.g., RAID1, RAID-5 3+1, RAIDS 7+1, and the like). Performance characteristics may relate to different performance aspects of the physical storage devices of a particular type or technology. For example, there may be multiple types of FC disk drives based on the RPM characteristics of the FC disk drives (e.g., 10K RPM FC drives and 15K RPM FC drives) and FC disk drives having different RPM characteristics may be included in different storage tiers. Storage capacity may specify the amount of data, such as in bytes, that may be stored on the drives. An embodiment may allow a user to define one or more such storage tiers. For example, an embodiment in accordance with techniques herein may define two storage tiers including a first tier of all SSD drives and a second tier of all non-SSD drives. As another example, an embodiment in accordance with techniques herein may define three storage tiers including a first tier of all SSD drives which are flash drives, a second tier of all FC drives, and a third tier of all SATA drives. The foregoing are some examples of tier definitions and other tier definitions may be specified in accordance with techniques herein.

Referring to FIG. 3, shown is an example **100** of software that may be included in a service processor such as **22a**. It should be noted that the service processor may be any one of a variety of commercially available processors, such as an Intel-based processor, and the like. Although what is described herein shows details of software that may reside in the service processor **22a**, all or portions of the illustrated components may also reside elsewhere such as, for example, on any of the host systems **14a 14n**.

Included in the service processor **22a** is performance data monitoring software **134** which gathers performance data

about the data storage system **12** through the connection **132**. The performance data monitoring software **134** gathers and stores performance data and forwards this to the optimizer **138** which further stores the data in the performance data file **136**. This performance data **136** may also serve as an input to the optimizer **138** which attempts to enhance the performance of I/O operations, such as those I/O operations associated with data storage devices **16a-16n** of the system **12**. The optimizer **138** may take into consideration various types of parameters and performance data **136** in an attempt to optimize particular metrics associated with performance of the data storage system **12**. The performance data **136** may be used by the optimizer to determine metrics described and used in connection with techniques herein. The optimizer may access the performance data, for example, collected for a plurality of LVs when performing a data storage optimization. The performance data **136** may be used in determining a workload for one or more physical devices, logical devices or volumes (LVs) serving as data devices, thin devices (described in more detail elsewhere herein) or other virtually provisioned devices, portions of thin devices, and the like. The workload may also be a measurement or level of "how busy" a device is, for example, in terms of I/O operations (e.g., I/O throughput such as number of I/Os/second, response time (RT), and the like).

The response time for a storage device or volume may be based on a response time associated with the storage device or volume for a period of time. The response time may be based on read and write operations directed to the storage device or volume. Response time represents the amount of time it takes the storage system to complete an I/O request (e.g., a read or write request). Response time may be characterized as including two components: service time and wait time. Service time is the actual amount of time spent servicing or completing an I/O request after receiving the request from a host via an HA **21**, or after the storage system **12** generates the I/O request internally. The wait time is the amount of time the I/O request spends waiting in line or queue waiting for service (e.g., prior to executing the I/O operation).

It should be noted that the operations of read and write with respect to an LV, thin device, and the like, may be viewed as read and write requests or commands from the DA **23**, controller or other backend physical device interface. Thus, these operations may also be characterized as a number of operations with respect to the physical storage device (e.g., number of physical device reads, writes, and the like, based on physical device accesses). This is in contrast to observing or counting a number of particular type of I/O requests (e.g., reads or writes) as issued from the host and received by a front end component such as an HA **21**. To illustrate, a host read request may not result in a read request or command issued to the DA if there is a cache hit and the requested data is in cache. The host read request results in a read request or command issued to the DA **23** to retrieve data from the physical drive only if there is a read miss. Furthermore, when writing data of a received host I/O request to the physical device, the host write request may result in multiple reads and/or writes by the DA **23** in addition to writing out the host or user data of the request. For example, if the data storage system implements a RAID data protection technique, such as RAID-5, additional reads and writes may be performed such as in connection with writing out additional parity information for the user data. Thus, observed data gathered to determine workload, such as observed numbers of reads and writes, may refer to the read and write requests or commands performed by the DA. Such read and write commands may correspond, respectively, to

physical device accesses such as disk reads and writes that may result from a host I/O request received by an HA **21**.

The optimizer **138** may perform processing of the techniques herein set forth in following paragraphs to determine how to allocate or partition physical storage in a multi-tiered environment for use by multiple applications. The optimizer **138** may also perform other processing such as, for example, to determine what particular portions of thin devices to store on physical devices of a particular tier, evaluate when to migrate or move data between physical drives of different tiers, and the like. It should be noted that the optimizer **138** may generally represent one or more components that perform processing as described herein as well as one or more other optimizations and other processing that may be performed in an embodiment.

Described in following paragraphs are techniques that may be performed to determine promotion and demotion thresholds (described below in more detail) used in determining what data portions of thin devices to store on physical devices of a particular tier in a multi-tiered storage environment. Such data portions of a thin device may be automatically placed in a storage tier where the techniques herein have determined the storage tier is best to service that data in order to improve data storage system performance. The data portions may also be automatically relocated or migrated to a different storage tier as the work load and observed performance characteristics for the data portions change over time. In accordance with techniques herein, analysis of performance data for data portions of thin devices may be performed in order to determine whether particular data portions should have their data contents stored on physical devices located in a particular storage tier. The techniques herein may take into account how "busy" the data portions are in combination with defined capacity limits and defined performance limits (e.g., such as I/O throughput or I/Os per unit of time, response time, utilization, and the like) associated with a storage tier in order to evaluate which data to store on drives of the storage tier. The foregoing defined capacity limits and performance limits may be used as criteria to determine promotion and demotion thresholds based on projected or modeled I/O workload of a storage tier. Different sets of performance limits, also referred to as comfort performance zones or performance zones, may be evaluated in combination with capacity limits based on one or more overall performance metrics (e.g., average response time across all storage tiers for one or more storage groups) in order to select the promotion and demotion thresholds for the storage tiers.

Promotion may refer to movement of data from a first storage tier to a second storage tier where the second storage tier is characterized as having devices of higher performance than devices of the first storage tier. Demotion may refer generally to movement of data from a first storage tier to a second storage tier where the first storage tier is characterized as having devices of higher performance than devices of the second storage tier. As such, movement of data from a first tier of flash devices to a second tier of FC devices and/or SATA devices may be characterized as a demotion and movement of data from the foregoing second tier to the first tier a promotion. The promotion and demotion thresholds refer to thresholds used in connection with data movement.

As described in following paragraphs, one embodiment may use an allocation policy specifying an upper limit or maximum threshold of storage capacity for each of one or more tiers for use with an application. The partitioning of physical storage of the different storage tiers among the applications may be initially performed using techniques herein in accordance with the foregoing thresholds of the application's

allocation policy and other criteria. In accordance with techniques herein, an embodiment may determine amounts of the different storage tiers used to store an application's data, and thus the application's storage group, subject to the allocation policy and other criteria. Such criteria may also include one or more performance metrics indicating a workload of the application. For example, an embodiment may determine one or more performance metrics using collected or observed performance data for a plurality of different logical devices, and/or portions thereof, used by the application. Thus, the partitioning of the different storage tiers among multiple applications may also take into account the workload or how "busy" an application is.

As an example, the techniques herein may be described with reference to a storage environment having three storage tiers—a first tier of only flash drives in the data storage system, a second tier of only FC disk drives, and a third tier of only SATA disk drives. In terms of performance, the foregoing three tiers may be ranked from highest to lowest as follows: first, second, and then third. The lower the tier ranking, the lower the tier's performance characteristics (e.g., longer latency times, capable of less I/O throughput/second/GB (or other storage unit), and the like). Generally, different types of physical devices or physical drives have different types of characteristics. There are different reasons why one may want to use one storage tier and type of drive over another depending on criteria, goals and the current performance characteristics exhibited in connection with performing I/O operations. For example, flash drives of the first tier may be a best choice or candidate for storing data which may be characterized as I/O intensive or "busy" thereby experiencing a high rate of I/Os to frequently access the physical storage device containing the LV's data. However, flash drives tend to be expensive in terms of storage capacity. SATA drives may be a best choice or candidate for storing data of devices requiring a large storage capacity and which are not I/O intensive with respect to access and retrieval from the physical storage device. The second tier of FC disk drives may be characterized as "in between" flash drives and SATA drives in terms of cost/GB and I/O performance. Thus, in terms of relative performance characteristics, flash drives may be characterized as having higher performance than both FC and SATA disks, and FC disks may be characterized as having a higher performance than SATA.

Since flash drives of the first tier are the best suited for high throughput/sec/GB, processing may be performed to determine which of the devices, and portions thereof, are characterized as most I/O intensive and therefore may be good candidates to have their data stored on flash drives. Similarly, the second most I/O intensive devices, and portions thereof, may be good candidates to store on FC disk drives of the second tier and the least I/O intensive devices may be good candidates to store on SATA drives of the third tier. As such, workload for an application may be determined using some measure of I/O intensity, performance or activity (e.g., I/O throughput/second, percentage of read operation, percentage of write operations, response time, etc.) of each device used for the application's data. Some measure of workload may be used as a factor or criterion in combination with others described herein for determining what data portions are located on the physical storage devices of each of the different storage tiers.

FIG. 4 is a schematic illustration showing a storage system 150 that may be used in connection with an embodiment of the system described herein. The storage system 150 may include a storage array 124 having multiple directors 130-132 and multiple storage volumes (LVs, logical devices or VOL-

UMES 0-3) 110-113. Host applications 140-144 and/or other entities (e.g., other storage devices, SAN switches, etc.) request data writes and data reads to and from the storage array 124 that are facilitated using one or more of the directors 130-132. The storage array 124 may include similar features as that discussed above.

The volumes 110-113 may be provided in multiple storage tiers (TIERS 0-3) that may have different storage characteristics, such as speed, cost, reliability, availability, security and/or other characteristics. As described above, a tier may represent a set of storage resources, such as physical storage devices, residing in a storage platform. Examples of storage disks that may be used as storage resources within a storage array of a tier may include sets SATA disks, FC disks and/or EFDs, among other known types of storage devices.

According to various embodiments, each of the volumes 110-113 may be located in different storage tiers. Tiered storage provides that data may be initially allocated to a particular fast volume/tier, but a portion of the data that has not been used over a period of time (for example, three weeks) may be automatically moved to a slower (and perhaps less expensive) tier. For example, data that is expected to be used frequently, for example database indices, may be initially written directly to fast storage whereas data that is not expected to be accessed frequently, for example backup or archived data, may be initially written to slower storage. In an embodiment, the system described herein may be used in connection with a Fully Automated Storage Tiering (FAST) product produced by EMC Corporation of Hopkinton, Mass., that provides for the optimization of the use of different storage tiers including the ability to easily create and apply tiering policies (e.g., allocation policies, data movement policies including promotion and demotion thresholds, and the like) to transparently automate the control, placement, and movement of data within a storage system based on business needs. The techniques herein may be used to determine amounts or allocations of each storage tier used by each application based on capacity limits in combination with performance limits.

Referring to FIG. 5A, shown is a schematic diagram of the storage array 124 as including a plurality of data devices 61-67 communicating with directors 131-133. The data devices 61-67 may be implemented as logical devices like standard logical devices (also referred to as thick devices) provided in a Symmetrix® data storage device produced by EMC Corporation of Hopkinton, Mass., for example. In some embodiments, the data devices 61-67 may not be directly useable (visible) to hosts coupled to the storage array 124. Each of the data devices 61-67 may correspond to a portion (including a whole portion) of one or more of the disk drives 42-44 (or more generally physical devices). Thus, for example, the data device section 61 may correspond to the disk drive 42, may correspond to a portion of the disk drive 42, or may correspond to a portion of the disk drive 42 and a portion of the disk drive 43. The data devices 61-67 may be designated as corresponding to different classes, so that different ones of the data devices 61-67 correspond to different physical storage having different relative access speeds or RAID protection type (or some other relevant distinguishing characteristic or combination of characteristics), as further discussed elsewhere herein. Alternatively, in other embodiments that may be used in connection with the system described herein, instead of being separate devices, the data devices 61-67 may be sections of one data device.

As shown in FIG. 5B, the storage array 124 may also include a plurality of thin devices 71-74 that may be adapted for use in connection with the system described herein when

using thin provisioning. In a system using thin provisioning, the thin devices 71-74 may appear to a host coupled to the storage array 124 as one or more logical volumes (logical devices) containing contiguous blocks of data storage. Each of the thin devices 71-74 may contain pointers to some or all of the data devices 61-67 (or portions thereof). As described in more detail elsewhere herein, a thin device may be virtually provisioned in terms of its allocated physical storage in physical storage for a thin device presented to a host as having a particular capacity is allocated as needed rather than allocate physical storage for the entire thin device capacity upon creation of the thin device. As such, a thin device presented to the host as having a capacity with a corresponding LBA (logical block address) range may have portions of the LBA range for which storage is not allocated.

Referring to FIG. 5C, shown is a diagram 150 illustrating tables that are used to keep track of device information. A first table 152 corresponds to all of the devices used by a data storage system or by an element of a data storage system, such as an HA 21 and/or a DA 23. The table 152 includes a plurality of logical device (logical volume) entries 156-158 that correspond to all the logical devices used by the data storage system (or portion of the data storage system). The entries in the table 152 may include information for thin devices, for data devices (such as logical devices or volumes), for standard logical devices, for virtual devices, for BCV devices, and/or any or all other types of logical devices used in connection with the system described herein.

Each of the entries 156-158 of the table 152 correspond to another table that may contain information for one or more logical volumes, such as thin device logical volumes. For example, the entry 157 may correspond to a thin device table 162. The thin device table 162 may include a header 164 that contains overhead information, such as information identifying the corresponding thin device, information concerning the last used data device and/or other information including counter information, such as a counter that keeps track of used group entries (described below). The header information, or portions thereof, may be available globally to the data storage system.

The thin device table 162 may include one or more group elements 166-168, that contain information corresponding to a group of tracks on the data device. A group of tracks may include one or more tracks, the number of which may be configured as appropriate. In an embodiment herein, each group has sixteen tracks, although this number may be configurable.

One of the group elements 166-168 (for example, the group element 166) of the thin device table 162 may identify a particular one of the data devices 61-67 having a track table 172 that contains further information, such as a header 174 having overhead information and a plurality of entries 176-178 corresponding to each of the tracks of the particular one of the data devices 61-67. The information in each of the entries 176-178 may include a pointer (either direct or indirect) to the physical address on one of the physical disk drives of the data storage system that maps to the logical address(es) of the particular one of the data devices 61-67. Thus, the track table 162 may be used in connection with mapping logical addresses of the logical devices corresponding to the tables 152, 162, 172 to physical addresses on the disk drives or other physical devices of the data storage system.

The tables 152, 162, 172 may be stored in the global memory 25b of the data storage system. In addition, the tables corresponding to particular logical devices accessed by a particular host may be stored (cached) in local memory of the

corresponding one of the HA's. In addition, an RA and/or the DA's may also use and locally store (cache) portions of the tables 152, 162, 172.

Referring to FIG. 5D, shown is a schematic diagram illustrating a group element 166 of the thin device table 162 in connection with an embodiment of the system described herein. The group element 166 may include a plurality of entries 166a-166f. The entry 166a may provide group information, such as a group type that indicates whether there has been physical address space allocated for the group. The entry 166b may include information identifying one (or more) of the data devices 61-67 that correspond to the group (i.e., the one of the data devices 61-67 that contains pointers for physical data for the group). The entry 166c may include other identifying information for the one of the data devices 61-67, including a speed indicator that identifies, for example, if the data device is associated with a relatively fast access physical storage (disk drive) or a relatively slow access physical storage (disk drive). Other types of designations of data devices are possible (e.g., relatively expensive or inexpensive). The entry 166d may be a pointer to a head of the first allocated track for the one of the data devices 61-67 indicated by the data device ID entry 166b. Alternatively, the entry 166d may point to header information of the data device track table 172 immediately prior to the first allocated track. The entry 166e may identify a cylinder of a first allocated track for the one of the data devices 61-67 indicated by the data device ID entry 166b. The entry 166f may contain other information corresponding to the group element 166 and/or the corresponding thin device. In other embodiments, entries of the group table 166 may identify a range of cylinders of the thin device and a corresponding mapping to map cylinder/track identifiers for the thin device to tracks/cylinders of a corresponding data device. In an embodiment, the size of table element 166 may be eight bytes.

Accordingly, a thin device presents a logical storage space to one or more applications running on a host where different portions of the logical storage space may or may not have corresponding physical storage space associated therewith. However, the thin device is not mapped directly to physical storage space. Instead, portions of the thin storage device for which physical storage space exists are mapped to data devices, which are logical devices that map logical storage space of the data device to physical storage space on the disk drives or other physical storage devices. Thus, an access of the logical storage space of the thin device results in either a null pointer (or equivalent) indicating that no corresponding physical storage space has yet been allocated, or results in a reference to a data device which in turn references the underlying physical storage space.

Thin devices and thin provisioning are described in more detail in U.S. patent application Ser. No. 11/726,831, filed Mar. 23, 2007 (U.S. Patent App. Pub. No. 2009/0070541 A1), AUTOMATED INFORMATION LIFE-CYCLE MANAGEMENT WITH THIN PROVISIONING, Yochai, EMS-147US, and U.S. Pat. No. 7,949,637, Issued May 24, 2011, Storage Management for Fine Grained Tiered Storage with Thin Provisioning, to Burke, both of which are incorporated by reference herein.

As discussed elsewhere herein, the data devices 61-67 (and other logical devices) may be associated with physical storage areas (e.g., disk drives, tapes, solid state storage, etc.) having different characteristics. In various embodiments, the physical storage areas may include multiple tiers of storage in which each sub-tier of physical storage areas and/or disk drives may be ordered according to different characteristics and/or classes, such as speed, technology and/or cost. The

devices **61-67** may appear to a host coupled to the storage device **24** as a logical volume (logical device) containing a contiguous block of data storage, as discussed herein. Accordingly, each of the devices **61-67** may map to storage areas across multiple physical storage drives. The granularity at which the storage system described herein operates may be smaller than at the file level, for example potentially as small as a single byte, but more practically at the granularity of a single logical block or collection of sequential data blocks. A data block may be of any size including file system or database logical block size, physical block, track or cylinder and/or other size. Multiple data blocks may be substantially the same size or different sizes, such as different size data blocks for different storage volumes or different sized data blocks within a single storage volume.

In accordance with techniques herein, an embodiment may allow for locating all of the data of a single logical portion or entity in a same tier or in multiple different tiers depending on the logical data portion or entity. In an embodiment including thin devices, the techniques herein may be used where different portions of data of a single thin device may be located in different storage tiers. For example, a thin device may include two data portions and a first of these two data portions may be identified as a "hot spot" of high I/O activity (e.g., having a large number of I/O accesses such as reads and/or writes per unit of time) relative to the second of these two portions. As such, an embodiment in accordance with techniques herein may have added flexibility in that the first portion of data of the thin device may be located in a different higher performance storage tier than the second portion. For example, the first portion may be located in a tier comprising flash devices and the second portion may be located in a different tier of FC or SATA drives.

Referring to FIG. 6, shown is an example illustrating information that may be defined and used in connection with techniques herein. The example **200** includes multiple storage tiers **206, 208, and 210**, an allocation policy (AP) **204**, and storage group (SG) **202**. The SG **202** may include one or more thin devices (TDs), such as TD A **220** and TD B **222**, used by an application **230**. The application **230** may execute, for example, on one of the hosts of FIG. 1. The techniques herein may be used to determine how to partition physical storage of the multiple storage tiers **206, 208 and 210** for use in storing or locating the application's data, such as data of the TDs **220 and 222**. It should be noted that the particular number of tiers, TDs, and the like, should not be construed as a limitation. An SG may represent a logical grouping of TDs used by a single application although an SG may correspond to other logical groupings for different purposes. An SG may, for example, correspond to TDs used by multiple applications.

Each of **206, 208 and 210** may correspond to a tier definition as described elsewhere herein. Element **206** represents a first storage tier of flash drives having a tier capacity limit C1. Element **208** represents a first storage tier of FC drives having a tier capacity limit C2. Element **210** represents a first storage tier of SATA drives having a tier capacity limit C3. Each of C1, C2 and C3 may represent an available or maximum amount of storage capacity in the storage tier that may be physical available in the system. The AP **204** may be associated with one of more SGs such as SG **202**. The AP **204** specifies, for an associated SG **202**, a capacity upper limit or maximum threshold for one or more storage tiers. Each such limit may identify an upper bound regarding an amount of storage that may be allocated for use by the associated SG. The AP **204** may be associated with one or more of the storage tiers **206, 208 and 210** that may be defined in a multi-tier storage environment. The AP **204** in this example **200**

includes limit **204a** identifying a maximum or upper limit of storage for tier1, limit **204b** identifying a maximum or upper limit of storage for tier2, and limit **204c** identifying a maximum or upper limit of storage for tier3. The SG **202** may be based on an SG definition identifying **202a** the logical devices, such as TDs included in the SG.

In connection with techniques herein, the maximum limits **204a, 204b and 204c** each represent an upper bound of a storage capacity to which an associated SG is subjected to. The techniques herein may be used to partition less than the amount or capacity represented by such limits. An amount of physical storage of a tier allocated for use by an application is allowed to vary up to the tier limit as defined in the AP **204** in accordance with other criteria associated with the application such as, for example, varying application workload. The optimizer may vary the amount of storage in each tier used by an SG **202**, and thus an application, based on workload and possibly other criteria when performing a cost benefit analysis, where such amounts are subject to the limits of the SG's AP and also performance limits described in more detail elsewhere herein. At a second point in time, the workloads and possibly other criteria for the applications may change and the optimizer may repartition the storage capacity used by each application subject to the capacity limits of APs and performance limits.

Referring to FIG. 7, shown is an example which more generally illustrates different associations between SGs, APs and tiers in an embodiment in accordance with techniques herein. The example **350** illustrates that an embodiment may have multiple storage tiers (e.g., tiers 1-N), multiple APs (e.g., AP1-N), and multiple SGs (e.g., SG 1-M). Each AP may be associated with one or more of the storage tiers. Each AP may also be associated with different tiers than other APs. For example, APn is associated with Tier N but AP1 is not. For each tier associated with an AP, the AP may define a maximum capacity limit as described in connection with FIG. 6. Each AP may be associated with one or more SGs. For example SGs1-N may be associated with a same AP1, and SGs N+1 through M may be associated with a same APn.

With reference back to FIG. 6, each of the maximum capacity limits may have any one of a variety of different forms. For example, such limits may be expressed as a percentage or portion of tier total storage capacity (e.g., such as a percentage of C1, C2, or C3), as an integer indicating an amount or quantity of storage **410c** (e.g., indicating a number of bytes or other number of storage units), and the like.

Data used in connection with techniques herein, such as the performance data of FIG. 3 used in determining device and SG workloads, may be obtained through observation and monitoring actual performance. Data may also be determined in other suitable ways such as, for example, through simulation, estimation, and the like. Observed or collected data may be obtained as described in connection with FIG. 3 by monitoring and recording one or more aspects of I/O activity for each TD, and portions thereof. For example, for each TD, and/or portions thereof, an average number of reads occurring within a given time period may be determined, an average number of writes occurring within a given time period may be determined, an average number of read misses occurring within a given time period may be determined, and the like. It should be noted that the operations of read and write with respect to an TD may be viewed as read and write requests or commands from the DA, controller or other backend physical device interface. Thus, these are operations may also be characterized as a average number of operations with respect to the physical storage device (e.g., average number of physical device reads, writes, and the like, based on physical device

accesses). This is in contrast to observing or counting a number of particular type of I/O requests (e.g., reads or writes) as issued from the host and received by a front end component such as an FA. To illustrate, a host read request may not result in a read request or command issued to the DA if there is a cache hit and the requested data is in cache. The host read request results in a read request or command issued to the DA to retrieve data from the physical drive only if there is a read miss. Furthermore, when writing data of a received host I/O request to the physical device, the host write request may result in multiple reads and/or writes by the DA in addition to writing out the host or user data of the request. For example, if the data storage system implements a RAID data protection technique, such as RAID-5, additional reads and writes may be performed such as in connection with writing out additional parity information for the user data. Thus, observed data gathered to determine workload, such as observed numbers of reads and writes, may refer to the read and write requests or commands performed by the DA. Such read and write commands may correspond, respectively, to physical device accesses such as disk reads and writes that may result from a host I/O request received by an FA.

It should be noted that movement of data between tiers from a source tier to a target tier may include determining free or unused storage device locations within the target tier. In the event there is an insufficient amount of free or unused storage in the target tier, processing may also include displacing or relocating other data currently stored on a physical device of the target tier. An embodiment may perform movement of data to and/or from physical storage devices using any suitable technique. Also, any suitable technique may be used to determine a target storage device in the target tier where the data currently stored on the target is relocated or migrated to another physical device in the same or a different tier.

One embodiment in accordance with techniques herein may include multiple storage tiers including a first tier of flash devices and one or more other tiers of non-flash devices having lower performance characteristics than flash devices. The one or more other tiers may include, for example, one or more types of disk devices. The tiers may also include other types of SSDs besides flash devices.

As described above, a thin device (also referred to as a virtual provision device) is a device that represents a certain capacity having an associated address range. Storage may be allocated for thin devices in chunks or data portions of a particular size as needed rather than allocate all storage necessary for the thin device's entire capacity. Therefore, it may be the case that at any point in time, only a small number of portions or chunks of the thin device actually are allocated and consume physical storage on the back end (on physical disks, flash or other physical storage devices). A thin device may be constructed of chunks having a size that may vary with embodiment. For example, in one embodiment, a chunk may correspond to a group of 12 tracks (e.g., 12 tracks*64Kbytes/track=768Kbytes/chunk). As also noted with a thin device, the different chunks may reside on different data devices in one or more storage tiers. In one embodiment, as will be described below, a storage tier may consist of one or more storage pools. Each storage pool may include multiple LVs and their associated physical devices. With thin devices, a system in accordance with techniques herein has flexibility to relocate individual chunks as desired to different devices in the same as well as different pools or storage tiers. For example, a system may relocate a chunk from a flash storage pool to a SATA storage pool. In one embodiment using techniques herein, a thin device can be bound to a particular storage pool of a storage tier at a point in time so

that any chunks requiring allocation of additional storage, such as may occur when writing data to the thin device, result in allocating storage from this storage pool. Such binding may change over time for a thin device.

A thin device may contain thousands and even hundreds of thousands of such chunks. As such, tracking and managing performance data such as one or more performance statistics for each chunk, across all such chunks, for a storage group of thin devices can be cumbersome and consume an excessive amount of resources. Described in following paragraphs are techniques that may be used in connection with collecting performance data about thin devices where such information may be used to determine which chunks of thin devices are most active relative to others. Such evaluation may be performed in connection with determining promotion/demotion thresholds use in evaluating where to locate and/or move data of the different chunks with respect to the different storage tiers in a multi-storage tier environment. In connection with examples in following paragraphs, details such as having a single storage pool in each storage tier, a single storage group, and the like, are provided for purposes of illustration. Those of ordinary skill in the art will readily appreciate the more general applicability of techniques herein in other embodiments such as, for example, having a storage group include a plurality of storage pools, and the like.

Referring to FIG. 8A, shown is an example 700 illustrating use of a thin device in an embodiment in accordance with techniques herein. The example 700 includes three storage pools 712, 714 and 716 with each such pool representing a storage pool of a different storage tier. For example, pool 712 may represent a storage pool of tier A of flash storage devices, pool 714 may represent a storage pool of tier B of FC storage devices, and pool 716 may represent a storage pool of tier C of SATA storage devices. Each storage pool may include a plurality of logical devices and associated physical devices (or portions thereof) to which the logical devices are mapped. Element 702 represents the thin device address space or range including chunks which are mapped to different storage pools. For example, element 702a denotes a chunk C1 which is mapped to storage pool 712 and element 702b denotes a chunk C2 which is mapped to storage pool 714. Element 702 may be a representation for a first thin device which is included in a storage group of one or more thin devices.

It should be noted that although the example 700 illustrates only a single storage pool per storage tier, an embodiment may also have multiple storage pools per tier.

Referring to FIG. 8B, shown is an example representation of information that may be included in an allocation map in an embodiment in accordance with techniques herein. An allocation map may be used to identify the mapping for each thin device (TD) chunk (e.g. where each chunk is physically located). Element 760 represents an allocation map that may be maintained for each TD. In this example, element 760 represents information as may be maintained for a single TD although another allocation map may be similarly used and maintained for each other TD in a storage group. Element 760 may represent mapping information as illustrated in FIG. 8A such as in connection the mapping of 702 to different storage pool devices. The allocation map 760 may contain an entry for each chunk and identify which LV and associated physical storage is mapped to the chunk. For each entry or row of the map 760 corresponding to a chunk, a first column 760a, Chunk ID, denotes an identifier to uniquely identify the chunk of the TD, a second column 760b, indicates information about the LV and offset to which the chunk is mapped, and a third column storage pool 760c denotes the storage pool and tier including the LV of 760b. For example, entry 762 represents chunk C1 illustrated in FIG. 8A as 702a and entry 764 represents chunk C2 illustrated in FIG. 8A as 702b. It should be

noted that although not illustrated, the allocation map may include or otherwise use other tables and structures which identify a further mapping for each LV such as which physical device locations map to which LVs. This further mapping for each LV is described and illustrated elsewhere herein such as, for example, with reference back to FIG. 5B. Such information as illustrated and described in connection with FIG. 8B may be maintained for each thin device in an embodiment in accordance with techniques herein.

In connection with collecting statistics characterizing performance, workload and/or activity for a thin device, one approach may be to collect the information per chunk or, more generally, for the smallest level of granularity associated with allocation and deallocation of storage for a thin device. Such statistics may include, for example, a number of reads/unit of time, # writes/unit of time, a number of prefetches/unit of time, and the like. However, collecting such information at the smallest granularity level does not scale upward as number of chunks grows large such as for a single thin device which can have up to, for example 300,000 chunks.

Therefore, an embodiment in accordance with techniques herein may collect statistics on a grouping of “N” chunks also referred to as an extent, where N represents an integer number of chunks, $N > 0$. N may be, for example, 480 in one embodiment. Each extent may represent a consecutive range or portion of the thin device in terms of thin device locations (e.g., portion of the address space or range of the thin device). Note that the foregoing use of consecutive does not refer to physical storage locations on physical drives but rather refers to consecutive addresses with respect to a range of addresses of the thin device which are then mapped to physical device locations which may or may not be consecutive, may be on the same or different physical drives, and the like. For example, in one embodiment, an extent may be 480 chunks ($N=480$) having a size of 360 MBs (megabytes).

An extent may be further divided into sub extents, where each sub extent is a collection of M chunks. M may be, for example 10 in one embodiment. In one embodiment, the sub-extent size may correspond to the smallest granularity of data movement. In other words, the sub extent size represents the atomic unit or minimum amount of data that can be operated upon when performing a data movement such as between storage tiers.

Referring to FIG. 9, shown is an example illustrating partitioning of a thin device’s address space or range in an embodiment in accordance with techniques herein. The example 250 includes a thin device address space or range 252 which, as described elsewhere herein, includes chunks mapped to physical storage locations. The thin device address space or range 252 may be partitioned into one or more extents 254a-254n. Each of the extents 254a-254n may be further partitioned into sub-extents. Element 260 illustrates that extent X 254n may include sub extents 256a-256n. Although only detail is illustrated for extent 254n, each of the other extents of the thin device also include a same number of sub extents as illustrated for 254n. Each of the sub extents 256a-256n may represent a grouping of “M” chunks. Element 262 illustrates that sub extent 1 256a may include chunks 258a-258n. Although only detail is illustrated for sub extent 256a, each of the other sub extents 256b-256n also include a same number of “M” chunks as illustrated for 256a. Thus, each of the extents 254a-254n may represent an grouping of “N” chunks, where

An embodiment in accordance with techniques herein may collect statistics for each extent and also other information characterizing activity of each sub extent of a thin device. Statistics for each extent may be characterized as either long term or short term. Short term refers to statistics which may reflect performance, workload, and/or I/O activity of an extent with respect to a relatively short window of time. Thus, short term statistics may reflect recent extent activity for such a short time period. In contrast and relative to short term, long term refers to statistics reflecting performance, workload and/or I/O activity of an extent with respect to a longer period of time. Depending on the evaluation being performed, such as by the optimizer, it may be desirable to place greater weight on short term information than long term, or vice versa. Furthermore, the information maintained per sub extent may be used as needed once particular extents of interest have been identified.

Referring to FIG. 10, shown is an example of information that may be collected and used in connection each extent in an embodiment in accordance with techniques herein. The example 300 illustrates that short term information 302, long term information 304 and a sub extent activity bitmap 306 may be collected for each extent. The short term information 302 and long term information 304 may be used in connection with determining short term rates 320 and long term rates 330 for each extent. The statistics included in 302, 304, 320 and 330 may reflect activity with respect to the entire extent. The activity bitmap 306 is illustrated in further detail by element 307 as including an entry for each sub extent in the associated extent. Entries of 307 are denoted by A, B, C, and the like. Each of the entries of 307 represents aggregated or collective activity information for a corresponding sub extent denoted by the numeric identifiers 307a of 1, 2, 3, etc. Each entry of 307 may include one or more bits used to encode an activity level with respect to all chunks of a corresponding sub-extent. For example, the entry of 307 denoted as A represents an activity level for all chunks in sub extent 1. An embodiment may use any number of bits for each entry of the activity bitmap 306, 307. For example, in one embodiment, each entry of the activity bitmap may be 2 bits capable of representing any of 4 integer values—0, 1, 2, and 3.

As will be described in following paragraphs, the short term rates 320, long term rates 330 and sub extent activity bitmap 306 may be used in connection with a variety of different evaluations such as by the optimizer 138. Generally, the activity level information or data for an extent such as illustrated in FIG. 10 may be referred to as extent activity level information including one or more metrics indicating an activity level for the extent. The extent activity level information may comprise short term activity information (e.g., such as 302 and/or 320) and long term activity information (e.g., such as 304 and 330).

In one embodiment, the short term rates 320 for an extent may include a read miss rate 322, a write I/O rate 324 and a prefetch rate 326 for the extent. The long term rates 330 for an extent may include a read miss rate 332 (e.g., number of read misses/unit of time, where a read miss refers to a cache miss for a read), a write I/O rate 334 (e.g., number of writes/unit of time) and a prefetch rate 336 (e.g., number of prefetches/unit of time) for the extent. As known in the art, data may be prefetched from a physical device and placed in cache prior to reference or use with an I/O operation. For example, an embodiment may perform sequential stream I/O recognition processing to determine when consecutive portions of a thin device are being referenced. In this case, data of the sequential stream may be prefetched from the physical device and placed in cache prior to usage in connection with a subsequent

I/O operation. In connection with a portion of data at a first point in a sequential stream associated with a current I/O operation, data subsequent to the first point may be prefetched such as when obtaining the portion from a physical device in anticipation of future usage with subsequent I/Os. Each of the foregoing rates of 320 and 330 may be with respect to any unit of time, such as per second, per hour, and the like. In connection with describing elements 302 and 304 in more detail, what will be described is how an embodiment in accordance with techniques herein may determine the short term rates 320 and long term rates 330 using a decay function and decay coefficients.

In an embodiment in accordance with techniques herein, a decay coefficient may be characterized as a weighting factor given to previous activity information. The higher the coefficient, the greater the weight given to previous activity information for the extent. Thus, the adjusted activity level of an extent at a current time, "An", may be generally represented as a function of a current observed or actual activity level for the current time, "a_n", a decay coefficient, "r", and previous adjusted activity level for the previous time period or sampling period, "A_{n-1}". In connection with the foregoing, "A" may represent an adjusted activity level, "n" may denote the current time period or sampling period and "n-1" may denote the immediately prior or previous time period or sampling period at which the activity for the extent was determined. In other words, "a_n" is adjusted to take into account previous activity as represented by "A_{n-1}" and "An" represents the resulting adjusted value of "a_n". With respect to a statistic or metric such as a number or read misses, "a_n" and "An" may each represent an integer quantity or number of read misses within a current sampling period, "n". The foregoing may generally be represented as:

$$A_n = a_n + (r * A_{n-1}) \tag{EQUATION 2}$$

wherein

- a_n is the actual observed activity metric for the current or "nth" sampling period,
- "r" is a decay coefficient,
- "A_n" is the adjusted activity metric for the current or "nth" sampling period, and
- "A_{n-1}" is the adjusted activity metric from the previous or "n-1" sampling period.

Beginning with an initial time period or sampling period, denoted by i="0" (zero), the adjusted activity A0 may be initially that which is observed, a0. Subsequent observed or actual activity levels may be adjusted as described above. Generally, "a_i" may denote an actual or observed value obtained for an activity metric for a sampling period "i", where "i" is an integer greater than or equal to 0. "A_i" may similarly denote an adjusted activity metric (or adjusted value for "a_i") for a sampling period "i", where "i" is an integer greater than or equal to 0. Thus, for consecutive sample periods at which actual or observed activity metrics are obtained (as denoted by lower case "a_i"s), corresponding adjusted activity levels (e.g., "A" values) may be determined as follows:

- A0=a0 /* Adjusted activity level A0, at time=0 or initially */
- A1=a1+(r*A0) /* Adjusted activity level A1, at first sampling period, i=1 */
- A2=a2+(r*A1) /* Adjusted activity level A2, at second sampling period, i=2 */
- :

so on for subsequent sampling periods 3, 4, and the like, based on EQUATION 2.

In connection with EQUATION 2, 0<=r<1, where "r" is a decay coefficient or weight given to previous activity. Varying

"r" in EQUATION 2 results in accordingly varying the weight given to past or previous activity. If r=0, then no weight is given to previous or historic values. Thus, the closer "r" is to 0, the lesser weight given to previous activity. Similarly, the closer "r" is to 1, the greater the weight given to previous activity. In connection with determining an adjusted activity level, An, using EQUATION 2 for short term and long term, different decay coefficients may be selected. Generally "r" for short term is less than "r" used in connection with long term activity. For example, in one embodiment, "r" used in connection short term activity levels may be 50% or 0.50 or smaller. "r" used in connection with long term activity levels may be 80% or 0.80 or larger. The foregoing are exemplary values that may be selected for "r" in connection with short term and long term activity levels depending on the weight to be given to previous activity. In connection with short term activity, a decay coefficient may be selected in accordance with providing a relatively short term rate of decay for an activity level metric determined at a point in time. For example, a short term rate of decay may provide for a rate of decay for an activity level metric on the order of one or more hours (e.g., less than a day). In connection with long term activity, a decay coefficient may be selected in accordance with providing a relatively long term rate of decay for an activity level metric determined at a point in time. For example, a long term rate of decay may provide for a rate of decay for an activity level metric on the order of one or more days, a week, and the like. Thus, an activity metric at a first point in time may have a weighted or residual effect on an adjusted activity level determined at a later point in time in accordance with the selected decay coefficient indicating the rate of decay of the activity metric.

As mentioned above, EQUATION 2 results in a metric or count, such as a number of read misses, number of writes, or number or prefetches during a sample period. It may be desirable to also determine a rate with respect to a unit of time, such as per second, per hour, and the like, for each of the foregoing adjusted activity metrics An. A rate with respect to a unit of time for the adjusted activity level An may be represented as:

$$Ar = An * (1-r) / (1-r^n) \tag{EQUATION 3}$$

where

- Ar=the adjusted activity rate per unit of time;
- r=decay coefficient or weight as described above,
- n=denotes an "nth" sampling period as described above,
- An=adjusted activity level determined for a given sampling period "n" (e.g. using EQUATION 2 as described above).

Generally, the higher the decay coefficient, r, the slower the change in Ar as may be the desired case with long term Ar values. Thus an embodiment may select decay coefficients for use with long term and short term Ar values so that, when plotted with respect to time, long term Ar values generally have a smaller slope than that associated with short term Ar values.

Referring to FIG. 11, shown is an example graphically illustrating the general shape of curves for long term (LT) and short term (ST) values in an embodiment in accordance with techniques herein. The activity level values (Y-axis values) are plotted with respect to time (X-axis). The activity level values may be determined using EQUATIONS 2 and/or 3. Curve 402 may be produced using one of EQUATIONS 2 and 3 where a first value for the decay coefficient "r" is selected for ST usage. Curve 404 may be produced using one of EQUATIONS 2 and 3 where a second value for the decay coefficient "r" is selected for LT usage. The values selected

for “r” in connection with **402** and **404** may be relative so that the first value for “r” used with **402** is less than the second value for “r” used with **404**.

In one embodiment, each of the different A_n values determined using EQUATION 2 may be converted to a corresponding A_r value using EQUATION 3 when desired.

In connection with the foregoing, for example, with respect to a number of read misses, “ a_n ” represents the number of such operations that have occurred in a current sample period, n . For example, if a sample period=10 minutes so that statistics for an extent are collected and/or computed every 10 minutes, “ a_n ” represents the number of read misses that occurred in the last 10 minute sample period or time interval. A_{n-1} represents the previous or last A calculation (e.g., as determined using EQUATION 2) from the previous sample period, denoted “ $n-1$ ”.

With reference back to FIG. 10, an embodiment may collect short term information **302** as counter values indicating a count or number of each type of operation for a current time period or sampling period “ n ”. The following may represent different “ a_n ” values as included in the short term information **302** for an extent: read miss count (number of read misses for the extent during the sampling period), prefetch count (number of prefetches for the extent during the sampling period) and write count (number of writes for the extent during the sampling period).

The short term information **302** may also include storing previous A values as determined for the sampling period “ $n-1$ ” using EQUATION 2 above. For example, short term information **302** may also include storing three (3) previous adjusted activity level values or A values for sampling period “ $n-1$ ” using EQUATION 2 above for the read miss count, prefetch count and write count.

The short term information **302** may also include a timestamp value indicating the timestamp associated with the previous sampling period “ $n-1$ ”.

Using the above-mentioned short term information **302**, an embodiment may calculate updated short term rates **320** using EQUATION 3 for a sampling period “ n ” for a selected “ r ” as a short term decay coefficient. With each new sampling period, the short term information may be accordingly updated so that which is associated with sampling period “ n ” subsequently becomes associated with sampling period “ $n-1$ ”.

The long term information **304** may include long term rates or A_r values as determined using EQUATION 3 for a read miss rate (e.g., number of read misses/second), a prefetch rate (e.g., number of prefetches/second) and a write rate (e.g., number of writes/second). The long term information **304** may also include a time duration interval used for determining an adjusted A_r value for the current time or sampling period “ n ”. For example, the time duration interval may represent the amount of time for which statistics are collected and used in connection with long term A_r values. An embodiment may store a set of long term A_r values rather than calculate such A_r values on demand from other stored information as in the case above for short term rates **320** (e.g., where short term information **302** is stored and used to calculate short term rates **320** on demand). Thus, in such an embodiment, the long term rates **330** may be included the long term information **304** where such long term rates **330** may be updated with each sampling period. In one embodiment with the arrival of a new sampling period “ n ”, the long term information **304** may include A_r values for the foregoing statistics as determined using EQUATION 3 for a sampling period “ $n-1$ ”. These long term A_r values for “ $n-1$ ” may each be multiplied by the time duration interval to determine A_{n-1} ,

an adjusted metric for the long term time period. The foregoing A_{n-1} value may then be used with EQUATION 2 to determine A_n for the current sampling period “ n ” using a selected “ r ” as a long term decay coefficient. Using A_n , EQUATION 3 may then be used to obtain updated long term rates A_r values. With each new sampling period, the long term information may be accordingly updated so that which is associated with sampling period “ n ” subsequently becomes associated with sampling period “ $n-1$ ”.

With reference back to FIG. 10, described above is an activity bitmap **306** having an entry per sub extent where each such entry may indicate an aggregate or collective activity level with respect to all chunks of the associated sub-extent. The number of different activity level states that may be represented for each sub extent depends on the number of bits per entry of the activity bitmap. In one embodiment, each entry of the activity bitmap may be 2 bits as described above so that each entry may be an integer in the inclusive range of 0 . . . 3. Processing may be performed to decrement each entry having a non-zero value by 1 every predetermined time period, such as every 12 hours. Each time there is any I/O operation to a sub extent since the sub extent was located or moved to its current physical location, the sub extent’s entry in the activity bitmap **306** may be set to 3. Thus, each entry in the bitmap may represent activity level information for up to 3 of the predetermined 12 hour time periods. An embodiment may also have a different number of bits per entry to represent a larger number of predetermined time periods. Based on the foregoing, the lower the value of a bitmap entry for a sub extent, the longer the amount of time that has lapsed since the sub extent has had any I/O activity.

Referring to FIG. 12, shown is a flowchart of processing steps that may be performed in connection with each activity bitmap associated with an extent in an embodiment in accordance with techniques herein. The flowchart **500** summarizes processing described above where each bitmap for each extent may be traversed with the occurrence of a predetermined time interval, such as every 12 hours. At step **502**, a determination is made as to whether the next time interval has lapsed. If not, processing waits at step **502** until step **502** evaluates to yes and control proceeds to step **504**. At step **504**, I is initialized to the next entry in the bitmap. I represents a loop counter when traversing through the bitmap and denotes the bitmap entry currently selected for processing. At step **506**, a determination is made as to whether the entire bitmap has been processed. If step **506** evaluates to yes, control proceeds to step **502** until an amount of time again lapses equal to that of the time interval. If step **506** evaluates to no, control proceeds to step **508** where a determination is made as to whether the current bitmap entry (e.g. bitmap [I]) is zero. If so, control proceeds to step **504**. Otherwise, control proceeds to step **510** where the current bit map entry is decremented by one (1) and control proceeds to step **504** to process the next entry in the bitmap.

The activity bitmap may be used in connection with determining an activity level associated with each sub extent, the smallest amount of data that can be associated with a data movement operation to relocate data from one physical device to another. It should be noted that an embodiment may have functionality and capability to physically move data in units or amounts less than a sub extent. However, when performing processing to determine data movement candidates, such as by the optimizer, such processing may consider candidates for data movement which have a minimum size of a sub extent. That is, all data of the sub extent may be either moved or relocated as a complete unit, or remains in its current location. In connection with a sub extent when per-

forming a data movement, it may be that not all chunks of the sub extent are actually moved. For example, suppose a sub extent is 10 chunks and the sub extent is to be moved from a first storage tier, such as from SATA or FC, to a second storage tier, such as flash. It may be that 9/10 chunks of the sub extent are unallocated or already in flash storage with only 1 chunk stored in the first storage tier. In this case, processing only needs to actually move the single chunk from the first storage tier to flash since the remaining 9 chunks are either already in the flash tier or unallocated. With a sub extent, the amount of data actually moved may be at most the size of the sub extent but may be less depending on, for example, whether all chunks of the thin device sub extent are allocated (e.g., actually map to physical storage), depending on the current physical device upon which chunks of the sub extent are located prior to movement, and the like. It should be noted that chunks of a sub extent may be located in different storage tiers, for example, depending on where the chunk's data is stored such as at the time when written as well as the result of other conditions that may vary with embodiment.

As an example use of the activity bitmap is in connection with promotion and demotion. As an example use of the activity bitmap, the bitmap may be used to determine selective sub extents which exhibit the highest activity level such as those having counters=3 (e.g., "hot" or active areas of the extent). These sub extents may be candidates for promotion or data movement to a higher performing storage tier and may be given preference for such promotion and data movement over other sub extents having activity bitmap entries which are less than 3. In a similar manner, the activity bitmap may be used to identify the "coldest" or inactive sub extents. For example, sub extents having bit map entries=0 may be candidates for demotion to a lower performing storage tier.

In connection with promotion data movements, an embodiment may want to be responsive to a change in workload with respect to the short term. With demotion, an embodiment may not want to move data as quickly as with promotion and may also want to consider longer term workloads prior to moving such data to a lesser performing storage tier. With promotion, an embodiment may give greater weight to ST workload and activity data. With demotion, an embodiment may additionally consider LT workload and activity rather than just such ST information.

The information as described and illustrated in FIGS. 10-12 above may be used for a variety of different purposes and evaluations. For example, an embodiment may use one or more of the short term rates to identify one or more active extents based on such aggregated extent-level activity data. Subsequently, once an active extent is identified such as a candidate for promotion, the extent's activity bitmap may be examined to determine which sub extents are most active. Processing may be performed to selectively move some of the sub extents of the active extent (e.g., those with counters=3) to a higher performing storage tier.

As another example, the activity bitmaps of extents may be used to determine a promotion ranking used to identify which extent may be promoted prior to one or more other extents. To further illustrate, an embodiment may have two extents, both which are candidates for promotion. The two extents may exhibit similar activity levels based on aggregate extent-level information such as based on short term rates 320 for each extent. The extent having the lesser number of active sub extents may have a higher priority for movement than the other extent. For example, processing may be performed to count the number of non-zero bit map entries for each of the two extents. The extent having the lower count may have a higher priority than the other extent having a higher count. In

other words, the extents may be ranked or ordered for promotion based on a number or count of non-zero bit map entries. The extent having the lower count may be characterized as also exhibiting the greatest activity level density based on the foregoing counts of the activity bitmaps.

As another example in connection with demotion, an embodiment may use one or more of the short term rates 320 in combination with one or more of the long term rates 330 to identify one or more inactive extents based on such aggregated extent-level activity data. Subsequently, once an inactive extent is identified, the extent's activity bitmap may be examined to determine which sub extents are inactive and should be demoted rather than automatically demoting all sub extents of the inactive extent. Processing may be performed to selectively move some of the sub extents (e.g., those with counters=0, counters less than some threshold such as 1, and the like) to a lower performing storage tier.

One embodiment in accordance with techniques herein may include multiple storage tiers including a first tier of flash devices and one or more other tiers of non-flash devices having lower performance characteristics than flash devices. The one or more other tiers may include, for example, one or more types of disk devices. The tiers may also include other types of SSDs besides flash devices.

The different levels of activity information described herein as obtained at a thin device level, extent level, and sub extent level provide a hierarchical view for characterizing activity of different portions of thin devices. Activity information at higher device levels may be used to first identify devices which may be candidates for data movement, such as between storage tiers (e.g. for promotion and/or demotion). In connection with thin devices, once such a first device is identified, additional detail regarding the first device's activity as reflected in extent activity level information may be used to identify an extent of the first device as a candidate for data movement. Subsequently, the activity bitmap for the extent identified may then be used to determine one or more sub extents of the identified extent for data movement. The techniques herein may be used for collecting and tracking activity of thin devices. Use of the decay coefficients and equations for determining adjusted activity levels to account for previous activity levels provides an effective way of tracking workload and activity over time without having to keep a large database of historical statistics and metrics for long and short time periods.

In addition to the activity information described above for each extent and sub extent of a thin device, an embodiment may also track device level activity information for logical devices (e.g., thin devices, LVs, and the like) and physical devices in a data storage system as also noted. Additionally, an embodiment may track activity information for thin device pools. When a DA or other device interface services an I/O, the DA may not typically have any-knowledge regarding thin devices as may be known from the host's point of view. In connection with collecting data for use with techniques herein, each DA may be provided with additional mapping information regarding thin devices and where storage for the thin devices is allocated (e.g., such as described by the allocation map). The DA may use this information to determine what thin device (if any) is associated with a given back end I/O request. When the DA is servicing a back end I/O request, the DA may record information about the I/O including information about the thin device associated with the I/O request. Such additional information about the thin device may be used in order to perform statistics collection of activity data for the thin devices in accordance with techniques herein.

In connection with techniques in following paragraphs, the extent-based short term and long term statistics or metrics as described in FIG. 10 may be used in determining scores indicating the activity of extents. In one embodiment, the score may be a weighted value based on a combination of all six metrics 322, 324, 326, 332, 334 and 336 of FIG. 10 although an embodiment may generally use any metrics in determining such scores. In an embodiment herein, a promotion score for an extent may be represented in EQUATION 4 as:

$$\frac{(p1*s_rrm+p2*s_w+p3*s_p+p4*l_rrm+p5*l_w+p6*l_p)}{(\# \text{ Active Subext}+1)}$$

where s_rrm is the rate of short term random read misses (322), s_w is the rate of short term writes (324), s_p is the rate of short term pre-fetches (326), l_rrm is the rate of long term random read misses (332), l_w is the rate of long term writes (334), and l_p is the rate of long term pre-fetches. The coefficients $p1$ - $p6$ may be set as appropriate. It should be noted that “# Active Subext” represents the number of active sub-extents or subportions of an extent or other larger data portion for which the score is being determined. Examples of evaluating when a subextent or other subportion is active are described elsewhere herein. It should be noted that metrics used in connection with determining promotion and/or demotion score may take into account I/O size.

In an embodiment herein, the values used may be 12, 4, 4, 3, 1, and 1, respectively. Of course, different values may be used to emphasize or deemphasize different I/O characteristics in connection with determination of the promotion raw score. The demotion score for an extent may be represented in EQUATION 5 as:

$$\frac{(p4*s_rrm+p5*s_w+p6*s_p+p1*l_rrm+p2*l_w+p3*l_p)}{(\# \text{ Active Subext}+1)}$$

where s_rrm , s_w , $p1$, etc. are as set forth above.

As noted above in connection with the exemplary EQUATIONS 4 and 5 for computing, respectively, the promotion and demotion scores, the same set of coefficients may be used. Alternatively, an embodiment may, however, use a different set of coefficients for computing the promotion and demotion scores.

In a multi-tiered storage system as described herein, an application having its data stored on thin devices of a storage group may be allowed to use multiple tiers of storage. In order to be able to use the storage of the tiers efficiently and also move a minimal number of chunks between tiers, chunks which are the most active or “hot” need to be located in the higher tiers (e.g., promoted to such tiers if not already located there) and chunks which are least active or “cold” need to be located in lower storage tiers (e.g., demoted to such tiers if not already located there). After identifying the hot and cold chunks, processing may be performed to determine how much of the hot chunks should be placed in the different storage tiers in order to efficiently utilize the higher performing tiers, such as flash tiers, while also avoiding overloading any given tier with I/O request or I/O transfer activity to the point that overall performance (e.g., across all tiers in the AP, across one or more SGs, for the whole data storage system, and the like with respect to the physical devices under consideration) would have been better had less of the workload been placed in the tier. In connection with the foregoing, techniques are described in following paragraphs which determine promotion and demotion thresholds of a data movement policy that may be associated with one or more SGs. The data movement policy as described herein in the context of thin devices affects what data portions of thin

devices are data movement candidates and may be moved to another tier. The selection of promotion and demotion thresholds may be made by considering criteria including performance limits (e.g., response time, number of I/Os per time period, and the like) and capacity limits. The performance limits may be flexible or adaptable and specified for each storage tier. The capacity limits may also be specified for each storage tier and may include capacity limits included in an AP for the affected one or more SGs. The techniques model response time of target storage tiers when evaluating different alternative hypothetical considerations in which performance limits are varied for each tier when selecting promotion and demotion thresholds. The different sets of performance limits in combination with capacity limits are evaluated by modeling the expected target tier performance and then determining an overall performance metric representing an aggregate modeled performance metric across all target storage tiers for all affected SGs. In one embodiment, the overall performance metric may be an average response time determined with respect to all target storage tiers using the modeled response time as determined for each such tier. The average response time is used to compare the overall modeled performance for the storage tiers when evaluating different sets of performance limits for each target tier. Each set of performance limits specified for multiple tiers may be used as a way to provide weighting factors for I/O workload distribution across the tiers in order to reflect the performance differences of the different tier storage technologies. Utilizing such “what if” analysis to evaluate different sets of performance limits coupled with capacity limits provides for determining promotion and demotion thresholds that may be used by the DA, or more generally, other backend data storage system components, in connection with performing data movements in accordance with workload or performance impact across all target storage tiers to increase overall performance.

In connection with techniques herein as mentioned above, response time may be considered as performance criteria alone, or in combination with other performance criteria in combination with capacity limits, when determining promotion and demotion thresholds affected what data portions of a thin device may be moved between physical storage devices in different storage tiers. The techniques herein consider different performance characteristic information and curves that may vary with each storage tier, type of physical device, device vendor, and the like. In particular, performance curves for the different storage tiers may be determined and used to model target tier and also overall SG performance across storage tiers as part of processing to evaluate different sets of performance limits in combination with capacity limits. As an example, consider a workload of N I/O operations/second. The response time experienced for the same workload varies with storage tier due to the underlying capabilities of each tier’s technology. As such, performance curves may be used in connection with techniques herein to model expected response times if a particular data movement is performed in accordance with candidate promotion and demotion thresholds.

Referring to FIG. 13, shown is an example of performance characteristic information illustrated in the form of curves for different storage tiers such as may be based on different disk drive types. The example 550 illustrates general curve shapes as may be associated with a SATA drive (as represented by 552) and an FC disk drive (as represented by 554) in connection with processing rate (X-axis in terms of I/Os/second) vs. response time (Y-axis). As may be seen from the illustration 550, for a same processing rate of I/Os/second, different RTs are obtained for each of a SATA drive and an FC disk drive. As

such, when moving data storage tier of SATA drives to a storage tier of FC drives, differences in performance characteristics such as response times are taken into consideration in accordance with techniques herein. An embodiment may store data as represented by the curves of FIG. 13 in one or more tables having rows and columns of data point values (e.g., X and Y coordinates for a plurality of points). When stored in tabular form, interpolation, curve fitting techniques, and the like, may be used in connection with determining values of X and Y coordinates lying between two existing points stored in the table. When considering moving data between devices of different types or more generally having different device characteristics, such tables of performance characteristic information may be used to determine, for a given processing rate of I/Os per second, a modeled RT for each of the different device types. For example, consider a first storage tier of SATA drives and a second storage tier of FC disk drives. In modeling performance based on a proposed data movement, an aggregated or total processing rate for each target tier may be determined, for example, using performance data collected. For such a total processing rate on the X-axis, a corresponding modeled RT value (Y-axis) may be obtained for each storage tier using tables or curves, such as illustrated in FIG. 13. An embodiment may use appropriate performance curves for each of the different storage tiers and associated technologies of the tiers. The performance curves may be obtained for each storage tier based on observed or collected data through experimentation. The particular parameters or metrics of collected data used to obtain performance curves to model expected RT may vary with storage tier and underlying technology. For example, as described in U.S. patent application Ser. No. 12/924,361, filed Sep. 24, 2010, TECHNIQUES FOR MODELING DISK PERFORMANCE, which is incorporated by reference herein, performance curves for modeling response times for disk drives is described using total number of I/Os and I/O size. Other technologies such as flash-based drives may use other parameters in modeling to determine the appropriate performance curve. For example, one approach to modeling flash-based drives may utilize observed performance data related to total number of I/Os, I/O size, and a ratio of read operations/write operations. Additionally, data modeling for different storage drives may utilize a feedback process. At a point in time, there is a set of data representing the performance curve for a particular drive. The actual measured RT of the drive for a given workload in terms of I/Os per second, for example, may be compared to a modeled RT value determined using the performance curve for similar model parameter values. Adjustments may be made to the modeled performance curve based on differences between the measured RT and modeled RT.

In connection with estimating thin device workloads, various metrics that may be used are described herein and also in U.S. patent application Ser. No. 12/924,396, filed Sep. 25, 2010, TECHNIQUES FOR STATISTICS COLLECTION IN CONNECTION WITH DATA STORAGE PERFORMANCE, which is incorporated by reference herein. Workload for thin devices may be determined in a variety of different ways in connection with determining the contributions of the thin device data portions that may be stored in multiple thin device pools. One approach may be to examine the allocation map and determine the workload of data portions based on I/Os directed to the physical device where such data portions are stored. However, an embodiment may use alternative approaches to estimate thin device workload due to additional resources consumed in connection with use of the allocation map which may adversely impact performance.

When data portions of a thin device are moved from a first storage tier to a second storage tier, the related workload of such data portions are moved to the target tier. In one embodiment, storage for thin devices may be evenly distributed across a pool of data devices comprising a thin device pool. This results in even distribution of capacity and I/O workload thereby making it possible to correlate I/O workload and capacity allocation at the pool level rather than reading the allocation map for each thin device. In other words, a workload for a thin device data portion having storage allocated from a thin device pool of data devices may be estimated by collecting thin device pool statistics and then apportioning an amount of the workload indicated by the collected data distributed evenly across all data portions stored in the pool.

As discussed elsewhere herein, policies may be used to determine when to promote data (map the data to a relatively faster tier) and when to demote data (map the data to a relatively slower tier). In particular, one such policy is a data movement policy based on promotion and demotion thresholds that may be determined using promotion and demotion scores for data portions. In an embodiment herein, this may be performed by first determining a score for different portions of a storage space based on relative activity level and then constructing promotion and demotion histograms based on the different scores and the frequency of each. In connection with thin devices, each of the data portions may correspond to a logical extent for which such scores are determined. Exemplary ways in which the promotion and demotion scores may be calculated are described above. The promotion and demotion scores may be used, respectively, in connection with the promotion and demotion histograms described below in more detail. Generally, the scores may be characterized as reflecting the I/O benefit to the host application and cost (e.g., in terms of performance bandwidth) to the targeted storage device tier. In connection with constructing the histogram, all extents are ordered or sorted according to their scores, from highest to lowest. Those extents having the highest scores are generally those preferred to be selected for having storage allocated from the highest performing tier. The histogram is one way in which such scores may be sorted and utilized in connection with techniques herein. It will be appreciated by those of ordinary skill in the art that there are alternative ways to define and compute the scores than as described herein. In one embodiment described herein, the scores may be computed differently for promotion and demotion to reflect the difference in criteria related to data movement into and out of storage tiers.

For purposes of illustration, consider an example of a single SG which may use a group of data devices, and thus physical devices, in three thin device pools—one for each of three storage tiers such as illustrated in FIG. 8A. Workload statistics such as described in connection with FIG. 10 may be computed for each extent and a promotion score may be calculated for each extent in the SG. Also, assume that only thin devices managed in accordance with techniques herein for which data movement may be performed are located in the SG and use the foregoing thin device pools.

Referring to FIG. 14, a histogram 1000 illustrates a plurality of activity bins (buckets) and the frequency thereof. Each vertical line of the histogram 1000 represents a bin corresponding to a number of data portions (e.g., extents) having the corresponding score. Determination of a score for a data portion is discussed in more detail elsewhere herein. In an embodiment herein, there are five thousand bins. Of course, a different number of bins may be used instead. The height of each bin represents a number (frequency) of data portions having a particular score. Thus, the longer a particular vertical

line, the more data portions there are having the corresponding score. Note that the sum of all of the frequencies of the histogram equals the total number of data portions of the system. Note also that the sum of frequencies of a portion between a first score and a second score equals the total number of data portions having a score between the first and second scores. As such, the total capacity allocated for a particular bin assuming a fixed size data portion may be determined as the mathematical product of the frequency of data portions in the bin (of those data portions having allocated storage) and the size of a data portion. If the data portions in a bin may have varying size, then such sizes corresponding to the allocated storage amounts for the data portions may be summed to determine the total capacity of storage allocated for the bin. In a similar manner, the modeled response time (e.g., average) for the total cumulative workload (e.g., total I/Os/second) of those data portions may be determined. The histogram 1000 also shows a first range indicator 1002 and a second range indicator 1004. The first range indicator 1002 corresponds to bins having a score from S1 to SMAX (the maximum score). The second range indicator 1004 corresponds to bins having a score of S2 to S1-1. In an embodiment herein, there are three levels or tiers of physical storage and data portions of the thin device having a score corresponding to the first range indicator 1002 are promoted (mapped) to a highest (fastest) level of storage and data portions having a score corresponding to the second range indicator 1004 are promoted (mapped) to a medium level of storage. The remainder of the portions (i.e., that do not correspond to either the first range indicator 1002 or the second range indicator 1004) are not changed (are not promoted based on the histogram 1000). Of course, it is possible to have any number of storage levels. Thus, S1 may represent the promotion score corresponding to the promotion threshold for the first or highest storage tier so that all data portions having a score at or above S1 are promoted to the highest storage tier, or otherwise considered a candidate for such promotion if not already located in the highest storage tier. In a similar manner, S2 may represent the promotion score corresponding to the promotion threshold for the second storage tier (next fastest) so that all data portions having a score at or above S2 and below S1 are promoted to the second storage tier, or otherwise considered a candidate for such promotion if not already located in the second storage tier. It should be noted that histogram of FIG. 14 used in connection with promotion scores and also FIG. 16 (described below) in connection with demotion scores may include scores for all data portions under consideration or analysis. For example, as described elsewhere herein in connection with other examples, the techniques herein may be performed with respect to a number of storage groups of thin devices having their storage allocated from one or more storage pools so that the thin devices have storage allocated from a set of physical drives. In this case, the histograms may include scores with respect to the foregoing data portions of the number of storage groups under consideration and evaluation with the techniques herein.

It should be noted that an embodiment using a histogram may select a suitable number of bins or buckets and an interval for each such bin. In one embodiment, the size of each bin may be driven by a selected number of bins with each bin having the same size. Additionally, an embodiment may use different techniques in connection with mapping or converting the promotion and demotion scores to indices associated with histogram bins. For example, an embodiment may use linear scaling to set a lower boundary for buckets having an associated index lower than a selected pivot value and may

use logarithmic scaling to set a lower boundary for buckets above the pivot. Logarithmic scaling may be appropriate in embodiments having larger scores or a wide range of scores in order to scale the size of scores above the pivot. In such embodiments, the score range associated with a bucket interval above the pivot varies so that a reasonable number of data portions are mapped to the associated bucket. Whether a histogram or other suitable technique is used may vary with the number of buckets, the number of data portions, and the like.

Additionally, it should be noted that rather than have a histogram with frequency on the Y-axis as in FIG. 14, an embodiment may represent the total allocated capacity on the Y-axis of the number of data portions having scores within a particular bin. In other words, the height of the bucket or bin represents the total allocated capacity of the scores mapped to that bin. Other representations are possible besides histograms in connection with determining promotion thresholds and also demotion thresholds as described elsewhere herein in more detail.

In connection with determining the first tier promotion threshold S1 of FIG. 14, processing is performed to map a number of data portions to the highest performing tier in accordance with criteria including a combination of one or more capacity limits and one or more performance limits. A capacity limit may be specified for each storage tier for the SG in an AP associated with the SG as described above. Additionally, a capacity limit indicating the physical maximum amount of storage capacity as a physical characteristic of the drives may also be considered since it may be possible in some embodiment to exceed the maximum capacity of the drives prior to exceeding the capacity limits in accordance with an AP. Additionally, one or more sets of performance limits may be specified for each storage tier. In one embodiment, performance limits may be specified in terms of response time for each tier. An embodiment may define one or more sets of predetermined response time performance limits for storage tiers where such sets of response time limits may also referred to as performance or comfort zones. Each set contains a response time limit for each storage tier that may be the target of promotion. In one embodiment, limits are not specified for the bottom tier. In one embodiment, seven comfort zones may be specified where each zone includes a response time limit for the first highest performing storage tier, such as flash-based tier, and the second highest performing tier, such as FC disk drives. For example, the following response time performance limits may be specified for seven comfort zones in the embodiment having 3 storage tiers:

Comfort Zone	EFD/flash Response Time (ms)	FC disk Response Time (ms)
1	1	6
2	2	10
3	3	14
4	4	18
5	6	25
6	8	40
7	10	50

Of course, an embodiment may provide any number of comfort zones more or less than seven and for a different number of storage tiers. Additionally, the foregoing values are exemplary and may vary with technology, drive vendor, and the like. Generally, values specified as the performance limit metrics, such as response times, may vary with the workload of a particular system and may be determined in any suitable

manner. For example, values for the foregoing metrics may be made based on knowledge regarding particular workload of a system and typical performance of drives of different storage tiers in a system. In this manner, limits specified may be realistic and in accordance with typical workload performance within a system. It should be noted that the foregoing limits may also be selected based on end user performance requirements. Additionally, as noted elsewhere herein, although response time is used as the workload or performance metric in connection with the foregoing comfort zones, other performance criteria metrics may be used in combination with, or as an alternative to, response time. For example, an embodiment may use utilization as a metric in a manner similar to response time in connection with techniques herein. That is, just as comfort zones include response time limits for storage tiers, comfort zones may include other criteria such as a utilization for each storage tier. As known in the art, utilization of a resource, such as a physical drive or with respect to physical drives of a storage tier, may be defined as a metric measuring an amount of time a device is utilized or in a non-idle state. For example, utilization for a storage tier may be represented as a percentage (e.g., based on a ratio of an amount of time the physical devices of the storage tier are in the non-idle state/total amount of time). The foregoing utilization metric may represent the average utilization for a storage tier determined over a period of time.

Generally, processing may be performed to determine a set of promotion thresholds S1, S2 of FIG. 14 in accordance with criteria including capacity limits and a set of performance limits for a single comfort zone. The processing traverses the histogram, from highest score to lowest score, mapping data portions to the first storage tier until either the capacity limit for the first storage tier is reached or until the response time performance limit for the first storage tier is reached. For each storage tier, a performance counter is maintained indicating a modeled current I/O processing rate and associated modeled response time based on those data portions currently mapped to the storage tier. As described elsewhere herein, performance curves such as illustrated in FIG. 13 may be used in modeling current performance for each storage tier based on data portions currently mapped to the storage tier when traversing the histogram scores. As each bucket or bin of the histogram has its data portions mapped to the first storage tier, the performance counter (indicating an updated modeled tier RT) is updated to reflect the modeled performance for the first storage tier as also including the additional data portions of the bucket now newly mapped to the first storage tier. For example, as a bucket of data portions is mapped to the first storage tier, the performance or workload information attributed to the newly added data portions in combination with those data portions already mapped to the first storage tier may be input to the appropriate storage tier performance model to determine a modeled aggregate response time. For example, as described above, one disk performance model for SATA and FC disk drives may use as the following as modeling inputs—total number of I/Os (e.g., used to determine the number of I/Os per second or other unit of time) and I/O size (or average I/O size of the total number of I/Os considered)—as collected or observed for the data portions. With these modeling inputs for the aggregated data portions mapped to the first storage tier, the modeling technique may use performance curves to determine an estimated or modeled response time for the physical storage devices in the storage tier based on the aggregate workload of the existing data portions currently mapped to the first storage tier and the additional data portions now also mapped to the first storage tier. In a similar manner, processing may track the current

amount of storage of the first tier consumed via the mapping so far. After each bucket of data portions is additionally mapped to the first storage tier to hypothetically represent or model movement of such data portions to the first storage tier, a determination may be made as to whether any of the capacity limits or the response time performance limit for the first tier has been reached or exceeded. If so, the score associated with the current bucket is the promotion threshold. Thus, all data portions in buckets higher than the current bucket (e.g., scores exceeding that of the current bucket) are candidates for promotion to the first storage tier. It should be noted that in connection with the foregoing promotion threshold, the score used as the promotion threshold may be the upper limit of the bucket interval (e.g., score range) for the current bucket at which at least one of the capacity limits or response time performance limits was exceeded during histogram traversal.

In connection with response time performance modeling for a storage tier, as described elsewhere herein with thin devices, the additional I/Os associated with the data portions being added (via mapping) to a storage pool of a particular storage tier may be modeled as being evenly distributed across drives of the storage pool. In the simplified example described herein with only a single storage pool, the modeled storage pool response time is also the modeled storage tier response time. In the event of multiple storage pools in a single tier where all such pools are used by the SG, an embodiment may choose to evenly distribute the added I/O operations across all drives of the storage pool. As described elsewhere herein, a simplifying assumption is that there are no other consumers of the storage tier capacities than those thin devices under device management using the techniques herein. In the event that there are other types of devices having associated data stored on the storage tiers, the amount of storage consumed and the workload of such device may be considered when determining whether capacity and performance limits have been reached. It should be noted that the even distribution modeling as described above may reflect that which is actually performed by the storage tiers and devices therein being evaluated in connection with thin device storage allocation. If an embodiment allocates thin device storage in a different manner, then such modeling should reflect that which is performed in the embodiment.

In a similar manner, a promotion threshold for the second storage tier is determined by performing processing as described above for the first tier with the difference that the processing is performed for the second storage tier until either the capacity limits or response time performance limit of the first zone are reached for the second storage tier. The foregoing capacity limits and response time performance limits vary with each storage tier. Processing that maps data portions to the second storage tier resumes with the last unmapped data portion from the previous storage tier processing (e.g., the bucket of data portions of the histogram causing a limit of the first storage tier to be reached). In this manner, data portions which were not mapped to first tier storage are automatically mapped to storage in the next highest tier and processing for the second storage tier thereby resumes with the last unmapped data portion. At the end of the second storage tier processing for the current zone, the second storage tier promotion threshold S2 is determined as illustrated in FIG. 14.

Referring to FIG. 15, shown is a flowchart of steps summarizing processing as described above in connection with determining a single promotion threshold for a single target tier using criteria including capacity limits and comfort zone response time limits for the target tier as specified in a single zone of performance limits. Thus, flowchart 1050 may be executed twice to determine, for the first zone, the two pro-

motion thresholds **S1** and **S2** of FIG. 14. At step **1052**, initialization processing is performed. Step **1052** includes initializing a variable, **AMT**, that keeps track of the amount of storage portions to zero. Step **1052** also includes initializing an index variable, **I**, to the maximum score (highest bin). In an embodiment herein, there are five thousand bins, so **I** would be set to five thousand at the step **1054**. Of course, other numbers of bins are also possible. Following step **1052** is step **1054** where **AMT** is incremented by **FREQ[I]**, the amount of data mapped to bin **I**. Following the step **1054** is step **1056** where an updated modeled tier **RT** is determined. At step **1058**, a determination is made as to whether any of the capacity limits and/or response time performance limit for the current tier have been exceeded. Step **1058** may include comparing the updated modeled tier **RT** to the response time performance limit for the current zone and current target promotion tier. Step **1058** may include comparing the current amount of capacity of the target tier consumed via the modeled mapping represented by **AMT** to the **AP** capacity limit. As described elsewhere herein, the total capacity consumed across one or more bins may be determined based on the cumulative frequencies of those bins and the amount of allocated storage of the data portions in the foregoing one or more bins. Step **1058** may include comparing the current amount of capacity of the target tier consumed via the modeled mapping represented by **AMT** to the **SG** capacity limit such as may be based on the physical drive capacity limits. If it is determined at the test step **1058** that none of the established limits have been exceeded, then control passes from the test step **1058** to a step **1062** where the index variable, **I**, is decremented. Following the step **1062**, control passes back to the step **1054** for another iteration. If any one or more of the foregoing limits are exceeded, step **1058** evaluates to yes and control proceeds to step **1064** where a score threshold is assigned the value of **I**. Data portions having a score of **I** or higher are promoted to the highest level of storage. Following the step **1064**, processing is complete.

The methodology for determining score values used to map data portions (indicating promotion candidates) to one or more intermediate storage levels may be similar to that described above in connection with the flow chart **1050**. In the case of intermediate storage levels though, the index variable **I** would be initialized to a score that is one less than the lowest score of the next highest storage level. For example, if storage portions having a score of 4500 to 5000 are assigned to the highest storage level, then the index variable, **I**, would be initialized to **4499** in connection with determining scores for an intermediate storage level just below the highest storage level.

Once promotion threshold processing has completed for the current zone, demotion threshold processing is performed as will now be described.

Referring to FIG. 16, shown is a histogram **1100** similar to the histogram **1000**, discussed above which illustrates a plurality of scores and the frequency thereof. The histogram **1100** may be used to determine which of the data portions (if any) may be demoted (e.g., mapped to relatively slower physical storage). In some embodiments, the histogram **1100** may be identical to the histogram **1000**. In other embodiments, the histogram **1100** may be different than the histogram **1000** because the scores for the histogram **1000** used for promotion may be different than the scores for the histogram **1100** used for demotion. Determination of promotion and demotion scores is discussed in more detail elsewhere herein. The histogram **1100** also shows a first range indicator **1102** and a second range indicator **1104**. The first range indicator **1102** corresponds to bins having a score from **SMIN** (the

minimum score) to **S1**. The second range indicator **1104** corresponds to bins having a score of **S1+1** to **S2**. In an embodiment herein, there are three levels of storage and storage portions having a score corresponding to the first range indicator **1102** may be demoted (mapped) to a lowest (slowest) third tier of physical storage and storage portions having a score corresponding to the second range indicator **1104** may be demoted (mapped) to a medium or second tier of storage. The remainder of the portions (i.e., that do not correspond to either the first range indicator **1102** or the second range indicator **1104**) are not changed (are not demoted based on the histogram **1100**). Of course, it is possible to have any number of storage levels.

In an embodiment, the demotion threshold for a tier may be determined in any suitable manner. For example, an embodiment may select a demotion threshold with respect to demoting a data portion from a storage tier based on the threshold score determined as the promotion threshold for the storage tier. The demotion threshold may be selected as a score that is the same or lower than the promotion threshold. For example, the demotion threshold may be determined using a constant factor by which the promotion threshold for the same storage tier is multiplied. (e.g. promotion threshold for a tier= $1.2 \times$ demotion threshold for a storage tier). The foregoing may introduce a stationary zone between the promotion and demotion thresholds for a tier where scores falling this stationary zone are neither promoted or demoted with respect to the storage tier. Introduction of the stationary zone may serve as one mechanism that may be included in an embodiment to limit thrashing with respect to repeatedly promoting and then demoting the same data portions having scores which border the promotion or demotion threshold for a storage tier. The demotion threshold may be selected so that it is always equal to or less than the storage capacity for the **SG** as may be specified in an associated **AP**.

In an embodiment herein, the processing performed for demoting data portions (extents) may be similar to processing described in connection with FIG. 15 with the difference that processing may be reversed so that, for example, the portions to be demoted to the lowest level of storage may be determined prior to higher storage tiers by initially beginning with setting **I** in step **1052** to **SMIN** and incremented in each iteration. In such an embodiment, storage capacity limits and/or performance limits may be utilized as may be provided in connection with an embodiment. For example, an embodiment may not provide performance limits for the lowest/slowest performing tier but may provide such limits for other tiers. In this case, an embodiment may determine demotion thresholds based on the criteria provided (e.g., if performance limits are not provided for the third storage tier (e.g., slowest) then only capacity limits may be used for the third storage tier.

In some embodiments, when a data or storage portion (e.g., an extent) is selected for promotion, only active subportions (e.g., subextents) are promoted while inactive subportions remain at their current storage level. In an embodiment herein, a subportion is considered active if it has been accessed in the previous $4\frac{1}{2}$ days and is considered inactive otherwise. Of course, other appropriate criteria may be used to deem subportions either active or inactive. In some embodiments, when a data portion (e.g., an extent) is selected for demotion, the entire storage portion may be demoted, irrespective of activity level of subportions. In addition, in some embodiments, appropriate mechanism(s) may be provided to reduce the amount of data that is demoted so that more data is maintained on relative faster physical storage devices. Each extent may be evaluated for promotion first as described above and then for demotion if it has not otherwise

qualified for promotion. If an extent does not qualify for promotion or demotion, then no data movement is modeled for the extent and subsequently the extent is also not a candidate for data movement with respect to a set of criteria (e.g., capacity limits and performance zone limits) currently being evaluating through modeling using techniques herein. It should be noted that an extent that qualifies for promotion may not then subsequently be a candidate for demotion. Thus, a candidate that qualifies first for promotion may then be removed as a possible demotion candidate.

After processing is performed for the first and second storage tiers to determine promotion and demotion thresholds using capacity limits and the first zone's performance limits, an overall performance metric for the SG using the physical drives of the storage tiers just processed is determined. In one embodiment, this performance metric may be the modeled average response time (RT) for the SG across all storage tiers just processed and may be represented in EQUATION 6 as:

$$\text{Average RT} = (1/\text{Total I/Os per second}) * \sum \text{ALL_TIERS} \\ (\text{RT of tier} * \text{I/O operations per second for the tier})$$

In EQUATION 6, "Total I/Os per second" is the total number or aggregate of I/Os per second across all physical devices of the SG, " $\sum \text{ALL_TIERS}$ " is the mathematical summation of the product represented by "(RT of tier * I/O operations per second for the tier)". It should be noted that the "RT of tier" may represent the average response time of physical devices in a particular tier. Additionally, EQUATION 6 may generally be determined with respect to all SGs and devices thereof currently being evaluated using the techniques herein. The foregoing Average RT may serve as an overall metric regarding performance of the entire SG across all storage tiers considered to determine whether the modeled performance using the response time limits for the first zone is preferable over other response time limits of another zone. The foregoing EQUATION 6 is a weighted average response time calculation that considers the number of I/Os with a given response time. Alternatively, an embodiment may compute an average RT including separate weightings related to technology type. It should be noted in connection with computing the average RT for the SG using EQUATION 6, the RT for each storage tier of the SG is utilized. This RT for each storage tier may be the last modeled RT computed during the histogram traversal as a result of performing promotion and demotion threshold determination and modeling the performance of such proposed data movement candidate data portions. It should be noted that if other criteria, such as utilization, are used in addition to or as an alternative to RT, then an embodiment may compute an overall or average metric across all storage tiers similar to as described above with EQUATION 6. For example, if zones of performance limits are defined for utilization limits for the storage tiers, then a metric for computing average utilization across all storage tiers of devices being evaluated may be used to represent the overall performance criteria used in selecting a set of performance limits in combination with capacity limits, and also the associated promotion/demotion thresholds.

In a similar manner as just described for the first set of performance limits of the first zone, processing is also performed for the next zone 2 (e.g., using the second set of performance limits). Thus, promotion thresholds and an average RT using EQUATION 6 are produced as a result of processing in accordance with capacity limits in combination with performance limits of each zone. After each zone is processed for candidate promotion and demotion thresholds, a determination may be made as to whether to stop further evaluating remaining zones. Such a determination may be

made by comparing a first value for the average RT determined using EQUATION 6 for a current zone with second value for the average RT determined using EQUATION 6 for the previously processed zone. For example, after determining promotion and demotion thresholds using zone 1 performance limits in combination with capacity limits (zone 1 scenario) and then zone 2 performance limits in combination with capacity limits (zone 2 scenario), the average RT associated with the zone 1 scenario may be compared to the average RT associated with the zone 2 scenario. If the average RT for zone 2 scenario does not indicate a sufficient or threshold level of improvement over the average RT for zone 1, then no further zones may be evaluated. An embodiment may define a threshold value that represents the minimum amount of improvement expected in order to continue evaluating further zone scenarios (e.g., determining promotion and demotion thresholds using capacity limits and performance limits for subsequently defined zones): An embodiment may determine a difference in metric values obtained for the average RT for the two zone scenarios to be compared. An improvement between zone scenarios may be determined if there is decrease in the average RT (e.g., lower average RT means better overall performance). This decrease may be larger than the threshold in order for a sufficient level of improvement to be determined. Alternatively, an embodiment may set the threshold value to zero so that any decrease in average RT between scenarios is considered sufficient improvement to proceed with evaluating further zone performance limits in combination with capacity limits.

It should be noted that if one of the capacity limits has been exceeded on a preceding iteration of processing for the prior zone, processing using subsequent zones stops. The processing described herein assumes that the lowest storage tier has sufficient capacity to accommodate storage for any data portions not mapped to the other storage tiers.

Referring to FIG. 17, shown is a flowchart 1200 of steps that may be performed in an embodiment in evaluating and modeling performance for different performance limits in combination with capacity limits in an embodiment in accordance with techniques herein. The steps of 1200 summarize processing described above. At step 1202, one or more histograms may be constructed. In step 1204, current zone is set to 1 in connection with commencing processing for the first zone's performance limits. At step 1206, promotion and demotion thresholds are determined in accordance with the capacity limits and performance limits of the current zone. Selection of such thresholds is followed by modeling proposed data movements and determining modeled RTs for all storage tiers for the one or more SGs. At step 1208, the modeled average RT is determined as an overall performance metric across all storage tiers for the one or more SGs. At step 1210, a determination is made as to whether the first zone is currently being processed. If so, control proceeds to step 1214. Otherwise, control proceeds to step 1211 where a determination is made as to whether there has been sufficient improvement with respect to the modeled average RT values for the current zone scenario and the previous zone scenario. If step 1212 evaluates to no, processing stops. If step 1212 evaluates to yes, control proceeds to step 1214 where a determination is made as to whether the capacity limit has been reached. Step 1214 may examine any one or more capacity limits defined such as, for example, capacity limits (e.g., per storage tier, overall SG capacity limits, and the like) as may be defined in an AP, physical limits of drive capacities, and the like. If any one of these capacity limits has been exceeded, step 1214 may evaluate to yes and processing may stop. If step 1214 evaluates to no, control proceeds to step 1216 to

increment current zone to the next zone. At step **1218**, a determination is made as to whether this is the last zone. If so, processing stops. Otherwise, control proceeds to step **1206**.

It should be noted that FIG. **17** illustrates only one particular way in which the performance limit criteria and capacity limit criteria may be used in connection with selecting promotion and/or demotion thresholds based on stopping criteria. An embodiment may vary the stopping criteria. For example, an embodiment may perform the foregoing evaluation of all zones of performance limits and capacity limit(s) and determine an average RT value across all storage tier using EQUATION 6, for each such zone, without consideration of the stopping criteria at steps **1212** and/or **1214** and then select the performance zone limits resulting in the best relative average RT across all storage tiers. As another variation, an embodiment may terminate processing and evaluation of subsequent performance zone limits upon finding a first such zone having performance limits that results in a modeled average RT that is above a defined threshold. Thus, an embodiment in accordance with techniques herein may vary the stopping criteria specified in connection with FIG. **17**.

Once processing as described in FIG. **17** is completed, the promotion and demotion thresholds associated with the zone having performance limits resulting in the minimum average RT may be selected for implementation in connection with actually performing the previously modeled data movements. This is described and summarized now with reference to FIG. **18**.

With reference to FIG. **18**, at step **1302**, performance zone limits are selected having the minimum associated average response time as modeled. It should be noted that if other performance criteria and associated limits, such as in connection with utilization limits described elsewhere herein, is utilized, step **1302** may include considering other overall performance metrics besides the average response time across all storage tiers. For example, an embodiment may also consider the overall average utilization across all storage tiers. If the embodiment utilizes more than one overall performance metric, then step **1302** may include evaluating the combination of the overall performance metrics. For example, an embodiment may weight each overall performance metric in connection with step **1302** to select a particular performance zone and associated limit criteria. At step **1304**, data movements (e.g., promotion and demotions for the multiple storage tiers) may be performed based on criteria including the promotion and demotion thresholds determined for the selected performance zone limits of step **1302**. In step **1306**, performance zones may be re-evaluated as needed using techniques described herein. Additionally, the response time limits of the performance zones may also be modified as needed to adjust for any workload changes in the system. In other words, as described elsewhere herein, the performance zones defined should set forth reasonable response time limits based on workload of the system being evaluated. The performance zones may set forth response time criteria that varies as the system workload may vary in order to appropriately and automatically adjust response time limits to accommodate for such variations in workload dynamically. It should be noted that the re-evaluation at step **1306** may be performed in response to an occurrence of any suitable event. For example, such re-evaluation may be performed periodically (e.g., upon the occurrence of a predefined time interval), in response to measured or observed system performance reaching a threshold level (e.g., when the measured or moni-

tored response time of the data storage system reaches a defined threshold level), in response to a user's manual selection, and the like.

For purposes of simplification, examples above considered a single SG. An embodiment may evaluate multiple SGs in combination if they share physical devices or defined pools of devices so that there is a dependency in that they utilize the same data storage resources. Additionally, there may be other consumers of the physical devices beside those under management of an optimizer or other component using the techniques herein for data movement. For example, there may be devices which not under management of such a component performing data movement using techniques herein for any one or more reasons. When considering the performance limits of storage tiers, an embodiment may determine a performance baseline associated with such devices representing the workload of such devices in the system since such devices may be viewed as having consumed or utilized a portion of the allowable performance limits. The performance baseline may be defined as disk utilization or a response time value that a physical storage device or drive would have if the drive only hosted data storage for devices that are not under management by a component using the techniques herein. In one embodiment this may include those portions of thin devices which may not be moved between physical devices such as between storage tiers. An embodiment may determine the baseline performance in any suitable manner for unmovable thin devices. For example, an embodiment may determine the data or thick devices included in a thin device pool servicing the thin device and obtain performance data for each such data device in the thin pool. There is an assumption that the embodiment provides for an distribution of workload within pool data devices. Performance data may be obtained for each moveable thin device using the thin device pool where such performance data indicates the thin device workload as distributed over data devices of the thin pool. For each such data device, the workload associated with unmovable thin devices may be determined by subtracting the distributed moveable thin device workload associated with the data device from the observed workload of the data device. In other words, for a data device, the workload of the data device attributable to the moveable thin device is subtracted from the total workload of the data device. The result of the foregoing is an estimate of the data device workload attributable to non-moveable thin device portions.

In connection with the defined performance or comfort zones described herein, it should be noted that such zones are determined for the particular resource or service that may be consumed or utilized. In a similar manner, zones may be defined and evaluated in connection with other resources or services which are consumed or utilized in the data storage system. For example, zones and performance modeling variations may be modeled in connection with varying the amount of cache where cache limits may be placed on data cached for particular thick or data devices, thin devices, and other entities which consume cache. As another example, zones of performance limits may be specified for varying performance limits related to one or more DAs that service physical data storage devices. In a similar manner as described herein for storage tiers of physical devices, different performance criteria may be specified in terms of performance zones of limits. For example, with respect to DAs, utilization may be used as a performance metric for which comfort zones are defined.

In connection with avoiding thrashing, described herein are several techniques that may be utilized such as related to using weighting of long term and short term metrics (e.g., FIG. **10**) and using a stationary zone between demotion and

promotion thresholds for a storage tier. An embodiment may use different techniques to avoid large changes in promotion and demotion thresholds selected and utilized in successive time periods. An embodiment may determine a running average with respect to promotion and/or demotion thresholds determined using the techniques herein and use the running average as the actual threshold when implementing data movements. The running average of promotion and/or demotion thresholds may be determined, for example, over a period of time, or using N previous threshold values. An embodiment may also increase the number of performance zones evaluated.

It should be noted that the criteria which is evaluated using techniques herein may include capacity limits and performance limits. The processing performed herein provides for adaptive tier overloading protection by allowing the system to automatically select from different sets or zones of performance limits as system workload changes. The particular performance limit criteria of response time specified for each tier in each zone is only an example of a performance limit criteria that may be used in an embodiment. For example, performance limit criteria may use one or more other metrics other than response time, such as I/O processing rate (e.g., number of I/Os/second), # reads/second, # writes/second, service time, queue waiting time or wait time, length and/or number of wait queues, and the like. These one or more other metrics may be used alone or in combination with response time limits. Furthermore an embodiment may associate a different weighting factor with each of the different metrics included in performance limits specified for a zone. The weights used for each of the different metric may vary with performance zone. Furthermore, the actual metrics may also vary with performance zone. For example, it may be that for a first zone, a particular response time limit is being evaluated and other performance limit criteria is also included for evaluation. This additional performance limit criteria (e.g., an additional metric) may not be considered in evaluation with other response time limits of other zones.

Furthermore, the particular overall metric of average response time used to select between evaluated performance zones may vary in an embodiment from what is described herein. For example, an embodiment may use a different metric other than average response time, or may use the average response time metric, alone or in combination with, other overall performance criteria to evaluate and select between performance zone limits. For example, as described elsewhere herein, an embodiment may also use utilization as the performance metric, alone or in combination with, response time. In such an embodiment, comfort zones of utilization values may be specified and an average utilization may be determined across all storage tiers in a manner similar to calculating and using average response time in EQUATION 6. Utilization may also be modeled in a manner similar to response time as described, for example, in connection with FIG. 13 (e.g., use modeled utilization curves with I/Os per second on the X-axis and utilization on the Y-axis as may be determined through observed and collected data).

It should be noted that although the techniques described herein are used with thin devices providing virtual storage provisioning, the techniques herein may also be used in connection with other types of devices such as those not providing virtual provisioning.

The techniques herein may be performed by executing code which is stored on any one or more different forms of computer-readable media. Computer-readable media may include different forms of volatile (e.g., RAM) and non-

volatile (e.g., ROM, flash memory, magnetic or optical disks, or tape) storage which may be removable or non-removable.

While the invention has been disclosed in connection with preferred embodiments shown and described in detail, their modifications and improvements thereon will become readily apparent to those skilled in the art. Accordingly, the spirit and scope of the present invention should be limited only by the following claims.

What is claimed is:

1. A method for evaluating data movement alternatives comprising:

receiving a set of criteria including capacity limits and performance limits, said performance limits including a plurality of sets of performance limits, each of said plurality of sets including a plurality of performance limits for a plurality of storage tiers, each of said plurality of performance limits of said each set specifying a threshold level of performance for a different one of the plurality of storage tiers;

performing first processing to evaluate a plurality of alternatives for use in data movement with respect to a set of logical devices having data stored on a set of physical storage devices, each of said plurality of alternatives including a different set of data movement criteria comprising said capacity limits and a different one of said plurality of sets of performance limits, said set of physical storage devices comprising at least a first physical device of one of said plurality of storage tiers and a second physical device of another one of said plurality of storage tiers, wherein, for said each alternative, said first processing includes modeling performance of a first of the plurality of storage tiers when hypothetically storing first data portions of the logical devices in the first storage tier, said first data portions not exceeding a first performance limit, for the first storage tier, from the different one of the plurality of sets of performance limits, said first performance limit specifying a threshold level of performance for the first data portions of the logical devices stored in the first storage tier;

determining a plurality of overall performance metrics, wherein each of the plurality of overall performance metrics denotes average modeled performance across said set of physical devices for an associated one of the plurality of alternatives including an associated one of the plurality of sets of performance limits; and selecting one of said plurality of sets of performance limits having an associated one of the plurality of overall performance metrics that is better than any other of said plurality of overall performance metrics.

2. The method of claim 1, wherein each of said plurality of storage tiers includes storage devices having any of different technology and different performance characteristics than other storage devices included in others of said plurality of storage tiers.

3. The method of claim 2, wherein a data storage system includes at least three storage tiers comprising said first storage tier being a highest performing storage tier including solid state storage devices, a second storage tier being a middle performing storage tier including disk drives, and a third storage tier being a lowest performing storage tier including disk drives, and wherein said plurality of storage tiers includes at least said first storage tier and said second storage tier.

4. The method of claim 1, wherein said capacity limits include a capacity limit for a defined storage group of said set of logical devices.

5. The method of claim 1, wherein said first processing includes determining a promotion threshold for each of said plurality of storage tiers.

6. The method of claim 1, wherein said first processing includes determining a demotion threshold for each of said plurality of storage tiers.

7. The method of claim 1, wherein said logical devices include one or more thin devices, each of said thin devices being a virtually provisioned device.

8. The method of claim 7, wherein each of said thin devices has logical address range representing a presented storage capacity of said each thin device, and wherein at least a portion of said logical address range is not mapped to physical storage indicating that physical storage is not allocated for said portion.

9. The method of claim 1, wherein said method is performed in connection with optimization processing to optimize data storage system performance.

10. The method of claim 1, wherein said set of physical storage devices is used to store data portions of said set of logical devices and a second set of one or more logical devices, wherein said set of logical devices is defined to allow data movement between different ones of said physical storage devices and said second set of logical devices is defined to not allow data movement between different ones of said physical storage devices.

11. The method of claim 10, farther comprising:
determining a performance baseline for said second set of logical devices, said performance baseline being defined as a measurement of workload of said second set of logical devices that is directed to said set of physical storage devices.

12. The method of claim 1, wherein each of said plurality of sets of performance limits includes any one or more of a response time limit and a utilization limit for each of said plurality of storage tiers.

13. The method of claim 1, wherein each of said overall performance metrics is an average response time modeled for said set of physical devices for an associated one of the plurality of alternatives.

14. The method of claim 13, wherein said capacity limits include a plurality of maximum values each identifying a maximum amount of one of the plurality of storage tiers available for consumption, and wherein said first processing includes evaluating a first of said plurality of alternatives and proceeding to evaluate a second of said plurality of alternatives if said capacity limits have not been exceeded.

15. The method of claim 14, wherein evaluation of the second alternative is performed if modeling of said first alternative indicates that said first alternative is determined to improve performance with respect to said set of physical devices by a threshold amount.

16. A computer readable medium comprising code stored thereon for evaluating data movement alternatives, the computer readable medium comprising code stored thereon that, when executed, performs a method comprising:

receiving a set of criteria including capacity limits and performance limits, said performance limits including a plurality of sets of performance limits, each of said plurality of sets including a plurality of performance limits for a plurality of storage tiers, each of said plural-

ity of performance limits of said each set specifying a threshold level of performance for a different one of the plurality of storage tiers;

performing first processing to evaluate a plurality of alternatives for use in data movement with respect to a set of logical devices having data stored on a set of physical storage devices, each of said plurality of alternatives including a different set of data movement criteria comprising said capacity limits and a different one of said plurality of sets of performance limits, said set of physical storage devices comprising at least a first physical device of one of said plurality of storage tiers and a second physical device of another one of said plurality of storage tiers, wherein, for said each alternative, said first processing includes modeling performance of a first of the plurality of storage tiers when hypothetically storing first data portions of the logical devices in the first storage tier, said first data portions not exceeding a first performance limit, for the first storage tier, from the different one of the plurality of sets of performance limits, said first performance limit specifying a threshold level of performance for the first data portions of the logical devices stored in the first storage tier;

determining a plurality of overall performance metrics, wherein each of the plurality of overall performance metrics denotes average modeled performance across said set of physical devices for an associated one of the plurality of alternatives including an associated one of the plurality of sets of performance limits; and

selecting one of said plurality of sets of performance limits having an associated one of the plurality of overall performance metrics that is better than any other of said plurality of overall performance metrics.

17. The computer readable medium of claim 16, wherein each of said plurality of storage tiers includes storage devices having any of different technology and different performance characteristics than other storage devices included in others of said plurality of storage tiers.

18. The computer readable medium of claim 17, wherein a data storage system includes at least three storage tiers comprising said first storage tier being a highest performing storage tier including solid state storage devices, a second storage tier being a middle performing storage tier including disk drives, and a third storage tier being a lowest performing storage tier including disk drives, and wherein said plurality of storage tiers includes at least said first storage tier and said second storage tier.

19. The method of claim 1, wherein each of said plurality of sets of performance limits specifies a different performance limit for each storage tier that is a target for data promotion and does not specify a performance limit for a storage tier that is not a target for data promotion.

20. The method of claim 1, wherein said set of physical devices includes physical devices of N storage tiers, N being an integer greater than 2, and wherein each of said plurality of sets of performance limits includes N-1 performance limits for a highest performing N-1 of the N storage tiers and each of said plurality of sets of performance limits does not include a performance limit for a lowest performing of the N storage tiers.

* * * * *